arXiv:1510.06809v2 [q-fin.EC] 22 Jun 2016

A Link between Sequential Semi-anonymous Nonatomic Games and their Large Finite Counterparts

Jian Yang

Department of Management Science and Information Systems

Business School, Rutgers University

Newark, NJ 07102

Email: jyang@business.rutgers.edu

July 2015; revised, March 2016; accepted, May 2016

Abstract

We show that obtainable equilibria of a multi-period nonatomic game can be used by players in its large finite counterparts to achieve near-equilibrium payoffs. Such equilibria in the form of random state-to-action rules are parsimonious in form and easy to execute, as they are both oblivious of past history and blind to other players' present states. Our transient results can be extended to a stationary case, where the finite multi-period games are special discounted stochastic games. In both nonatomic and finite games, players' states influence their payoffs along with actions they take; also, the random evolution of one particular player's state is driven by all players' states as well as actions. The finite games can model diverse situations such as dynamic price competition. But they are notoriously difficult to analyze. Our results thus suggest ways to tackle these problems approximately.

Keywords: Nonatomic Game; Convergence in Probability; ϵ-Equilibrium

MSC2000 Classification Codes: 60G50; 91A10; 91A13

1 Introduction

We show that an equilibrium of a random multi-period game involving a continuum of players can be used to achieve asymptotically equilibrium results for its large finite counterparts. The latter finite games can model competitive situations involving random and action-dependent evolution of players' states which in turn influence period-wise payoffs. Their complex natures make equilibria difficult to locate. In contrast, those for the former continuum-player game are simple in form and relatively easy to obtain. Therefore, a bridge between the two types of games can have broad practical implications.

The former continuum-player game can be termed more formally as a *sequential semi-anonymous nonatomic game* (SSNG). In it, a continuum of players interact with one another in multiple periods; also, each player's one-time payoff and random one-period state transition are both swayed by his own state and action, as well as the joint distribution of other players' states and actions. This is indeed the anonymous sequential game studied by Jovanovic and Rosenthal [16]. We use the name SSNG just to be consistent with the single-period nonatomic-game (NG) literature, where anonymity has been reserved for a more special case. An SSNG's finite counterpart is almost the same except that only a finite number of players are involved. This more realistic situation is much more difficult to handle.

In a few steps, we demonstrate the usefulness of an SSNG equilibrium in finite multi-period games. First, in a precise language, Theorem 1 describes the gradual retreat of randomness in finite games as the number n of players tends to $+\infty$. This paves way for Theorem 2, which states that an SSNG's conditional equilibria, in terms of random state-to-action rules, can be used by its large finite counterparts to reach asymptotically equilibrium payoffs on average. A further refinement of this result is achieved in Theorem 3. The above transient results can be extended to the stationary case involving discounted payoffs and infinite horizons; see Theorem 4. The conditional equilibria that facilitate our study are similar to well-understood distributional equilibria. Their existence is also directly verifiable.

One practical situation to which our results can be applied concerns dynamic price competition. Here, players may be firms producing one identical product type, states may be combinations of the firms' inventory levels and other static or dynamic characteristics such as unit costs, and actions may be unit prices the firms charge for the product. In every period, the random demand arriving to a firm is dependent on not only its own price but also prices charged by other competitors. The actual sales is further constrained by available inventory. So the player's one-time payoff is a function of both its own state (inventory and probably

also cost) and action (price) and of the distribution of others' actions (prices). Moreover, the firm's next-period inventory level depends on its current level, the random demand, and potentially an exogenously given production schedule. So the random single-period state transition is potentially a function of the same factors involved in the payoff.

It is a difficult task to predict or prescribe what inventory-dependent prices the firms will or should charge over a finite time horizon. This can be further complicated by diverse scenarios where firms have different degrees of knowledge on their competitors' inventory levels and/or costs. Our results, on the other hand, will reveal that the nonatomic counterpart SSNG is easier to tackle. Its equilibria can be plugged back to the actual finite-player situations, without regard to the particularities of the scenarios, and still make reasonably well predictions/prescriptions when the number of players is large enough. We are not equipped to answer how large is "large enough". But computational study done in a related pricing setting hinted that player numbers "in the tens" seem large enough; see, Yang and Xia [35].

In the remainder of the paper, Section 2 surveys the relevant literature. We then spend Sections 3 and 4 on essentials of SSNGs and finite games, respectively. In Section 5, we demonstrate the key result that state evolutions in large finite games will not veer too far away from their NG counterparts. Section 6 is devoted to the main transient result and Section 7 its detailed interpretation. This result is extended to the stationary case in Section 8. Implications of these results and existence of our kind of equilibria for SSNGs are shown in Section 9. We conclude the paper in Section 10.

2 Literature Survey

From early on, NGs have been used as easier-to-analyze proxies of real finite-player situations, such as in the study of perfect competition. Systematic research on NG started with Schmeidler [28]. He formulated a single-period semi-anonymous NG, wherein the joint distribution of other players' identities and actions may affect any given player's payoff. When the action space is finite, Schmeidler established the existence of pure equilibria when the game becomes anonymous, so that other players' influence on a game's outcome is channeled through the marginal action distribution alone. Mas-Colell [24] showed the existence of distributional equilibria in anonymous NGs with compact metric action spaces. The latter result was extended by Khan and Sun [20] to a case where players differ on how their preferences over actions are influenced by external action distributions. A survey of related

works up until the early 2000s was provided by Khan and Sun [23].

Much attention has been paid to the topic of pure-equilibrium existence. Khan and Sun [21] developed a purification scheme involving a countable compact metric action space. Khan and Sun [22] used non-standard measures on identity spaces and generalized Schmeidler's pure-equilibrium existence result for more general action spaces. Balder [4] established pure- and mixed-equilibrium existence results that may be regarded as generalizations of Schmeidler's corresponding results. Other notable works still include Yu and Zhang [36] and Balder [5]. On the other hand, Khan, Rath, and Sun [18] identified a certain limit to which Schmeidler's result can be extended. Recently, Khan et al. [19] took players' diverse bio-social traits into consideration and pinpointed saturation of the player-identity distribution as the key to existence of pure equilibria.

Links between NGs and their finite counterparts were covered in Green [13], Housman [15], Carmona [8], Kalai [17], Al-Najjar [3], and Yang [33]. For multi-period games without changing states, Green [12], Sabourian [27], and Al-Najjar and Smorodinsky [2] showed that equilibria for large games are nearly myopic.

SSNGs are both challenging and rewarding to analyze because in them, very realistically, individual states are subject to sways of players' own actions as well as their opponents' states and actions. Jovanovic and Rosenthal [16] established the existence of distributional equilibria for such games. This result was generalized by Bergin and Bernhardt [6] to cases involving aggregate shocks. In SSNGs' finite-player counterparts, however, randomness in state evolution will not go away. Besides, a player's ability to observe other players' states and actions might also affect his decision. Presented with these difficulties, it is not surprising that known results on sequential finite-player games are restricted to the stationary setting, where they appear as discounted stochastic games first introduced by Shapley [29]. According to Mertens and Parthasrathy [25], Duffie, Geanakoplos, Mas-Colell, and McLennan [9], and Solan [30], for instance, equilibria known to exist for these games come in quite complicated forms that for real implementation, demand a high degree of coordination among players.

It is therefore natural to ask whether sequential finite-player games can be approximated by their NG counterparts. This question has so far been answered by two unpublished articles. For a case unconcerned with the copula between marginal state and action distributions, Bodoh-Creed [7] provided an affirmative answer, and went on to show for certain cases that limits of large-game equilibria of a myopic form, when in existence, are NG equilibria. Also, Yang [34] verified for the approximability when both state transitions and action plans are

driven by exogenously generated idiosyncratic shocks. Our current study attempts with the most general possible setting, without unduly restricting on ways in which a player's payoff can be influenced by other players' states and actions or ways in which the game can evolve randomly. To achieve results of the same spirit, we have to overcome technical challenges posed by the new phenomenon of sampling from non-product joint probabilities.

Some authors went on to pursue stationary equilibria (SE), which stressed the long-run steady-state nature of individual action plans and system-wide multi-states; see, e.g., Hopenhayn [14] and Adlakha and Johari [1]. The oblivious equilibrium (OE) concept as proposed by Weintraub, Benkard, and van Roy [31], in order to account for impacts of large players, took the same stationary approach by letting participants beware of only long-run average system states. Weintraub, Benkard, and van Roy [32] showed links between equilibria of infinite-player games and their finite-player brethren for a setting where the long-run average system state could be defined. Though applicable to many situations, we caution that the implicit stationarity of SE or OE is incompatible with applications that are transient by nature; for instance, the dynamic pricing game studied by Yang [34].

3 The Nonatomic Game

The SSNG is a game in which a continuum of players interact with one another over multiple periods. A realistic and yet complicating feature is that players possess individual states which influence their payoffs along with all players' actions. The random evolutions of these states, meanwhile, are affected by players' actions. Furthermore, the semi-anonymous nature of the game means that not only what was done, but also who did what to the extent at which states partially reveal player identities, figure large in both payoff formation and state evolution. We now provide a detailed account of the game.

3.1 Game Primitives

For some natural number $\bar{t} \in \mathbb{N}$, we let periods $1, 2, ..., \bar{t}$ serve as regular periods and period $\bar{t}+1$ as the terminal period. For all periods, we let players' individual states and actions form, respectively, separable metric spaces S and X. We further require that both spaces be *discrete*. In this paper, such a space always stands for a separable metric space with countably many elements and the additional feature that the minimum of the distances between any two points remains strictly positive. The discreteness requirement will be useful

at one occasion. But most of our derivations will work if the spaces were merely separable metric. Given any separable metric space A, we use $\mathcal{B}(A)$ for its Borel σ-field and $\mathcal{P}(A)$ for the set of all probability measures on the measurable space $(A, \mathcal{B}(A))$.

To each player, other players' states and actions are immediately felt in a semi-anonymous fashion, so that what really matters is the joint distribution of other players' states and actions. This distribution, which we dub "in-action environment", is a member of the joint state-action distribution space $\mathcal{P}(S \times X)$. In any period $t = 1, 2, ..., \bar{t}$, a player's state $s \in S$, his action $x \in X$, and the in-action environment $\tau \in \mathcal{P}(S \times X)$ he faces, together determine his payoff in that period. In particular, there is a function

$$\tilde{f}_t : S \times X \times \mathcal{P}(S \times X) \to [-\bar{f}_t, \bar{f}_t], \tag{1}$$

where $\bar{f}_t$ is some positive constant on the real line $\mathbb{R}$. It is required that $\tilde{f}_t(\cdot, \cdot, \tau)$ be a measurable map from $S \times X$ to $[-\bar{f}_t, \bar{f}_t]$ for every $\tau \in \mathcal{P}(S \times X)$. For the terminal period $\bar{t} + 1$, we let the payoff be 0 in all circumstances.

Now we describe individual players' random state transitions. Given separable metric spaces A and B, we use $\mathcal{K}(A, B)$ to represent the space of all kernels from A to B. Each member $\kappa \in \mathcal{K}(A, B) \subseteq (\mathcal{P}(B))^A$ satisfies that

(i) $\kappa(a)$ is a member of $\mathcal{P}(B)$ for each $a \in A$, and

(ii) for each $B' \in \mathcal{B}(B)$, the real-valued function $\kappa(\cdot|B')$ is measurable.

Note that we have used $\kappa(a|B')$ rather than the more conventional $\kappa(B'|a)$ to denote the conditional probability for $B' \in \mathcal{B}(B)$ when given $a \in A$. The current notation allows us to always read a formula from left to right. Now in period $1, 2, ..., \bar{t}$, let there be a function

$$\tilde{g}_t : S \times X \times \mathcal{P}(S \times X) \to \mathcal{P}(S), \tag{2}$$

so that $\tilde{g}_t(\cdot, \cdot, \tau)$ is a member of $\mathcal{K}(S \times X, S)$ for each $\tau \in \mathcal{P}(S \times X)$. For convenience, we use $\mathcal{G}(S, X)$ to denote the space of all such functions, or what we shall call "state transition kernels". In period t, when a player is in individual state $s \in S$, takes action $x \in X$, and faces in-action environment $\tau \in \mathcal{P}(S \times X)$, there will be a $\tilde{g}_t(s, x, \tau|S')$ chance for his state in period $t + 1$ to be in any $S' \in \mathcal{B}(S)$.

This setup is versatile enough to embrace different player characteristics. For instance, each $s \in S$ may comprise two components θ and ω, with the $\tilde{g}_t$'s defined through (2) dictating that θ stays static over time to serve as a player's innate type. Certainly, the $\tilde{f}_t$'s defined through (1) can have all kinds of trends over θ to reflect players' varying payoff structures.

3.2 Evolution of the Environments

In any period $1, 2, ..., \bar{t}, \bar{t}+1$, by "pre-action environment" we mean the state distribution $\sigma \in \mathcal{P}(S)$ of all players. With $\bar{t}$, S, X, $(\tilde{f}_t | t = 1, 2, ..., \bar{t})$, and $(\tilde{g}_t | t = 1, 2, ..., \bar{t})$ all given in the background, we use $\Gamma(\sigma_1)$ to denote an (SS)NG with $\sigma_1 \in \mathcal{P}(S)$ as its initial period-1 pre-action environment. For this NG, we can use $\chi_{[1\bar{t}]} = (\chi_t \mid t = 1, ..., \bar{t}) \in (\mathcal{K}(S, X))^{\bar{t}}$ to denote a policy profile. Here, each $\chi_t \in \mathcal{K}(S, X)$ is a map from a player's state to the player's random action choice. Together with the given initial environment σ_1, this policy profile will help to generate a deterministic pre-action environment trajectory $\sigma_{[1,\bar{t}+1]} = (\sigma_t \mid t = 1, 2, ..., \bar{t}, \bar{t}+1) \in (\mathcal{P}(S))^{\bar{t}+1}$ in an iterative fashion. This process is also intertwined with the formation of in-action environments $\tau_1, \tau_2, ..., \tau_{\bar{t}}$ faced by all players in periods $1, 2, ..., \bar{t}$.

More notation is needed to precisely describe this evolution. Given distribution $p \in \mathcal{P}(A)$ and kernel $\kappa \in \mathcal{K}(A, B)$ for separable metric spaces A and B, there is a natural product $p \otimes \kappa \in \mathcal{P}(A \times B)$, such that

$$(p \otimes \kappa)(A' \times B') = \int_{A'} p(da) \cdot \kappa(a|B'), \qquad \forall A' \in \mathcal{B}(A), B' \in \mathcal{B}(B). \tag{3}$$

Here, $p \otimes \kappa$ is essentially the joint distribution generated by the marginal p and conditional distribution κ. Obviously, $(p \otimes \kappa)|_A$, the marginal of $p \otimes \kappa$ on A, is p. At the same time, we use $p \odot \kappa$ to denote the marginal $(p \otimes \kappa)|_B$, which satisfies

$$(p \odot \kappa)(B') = (p \otimes \kappa)|_B(B') = (p \otimes \kappa)(A \times B') = \int_A p(da) \cdot \kappa(a|B'), \qquad \forall B' \in \mathcal{B}(B). \tag{4}$$

Suppose pre-action environment $\sigma_t \in \mathcal{P}(S)$ has been given for some period $t = 1, ..., \bar{t}$. Then, for every player with starting state s_t in the period, his random action will be sampled from the distribution $\chi_t(s_t|\cdot)$ where as noted before, $\chi_t \in \mathcal{K}(S, X)$ is every player's behavioral guide. Thus, all players will together form the commonly felt in-action environment

$$\tau_t = \sigma_t \otimes \chi_t. \tag{5}$$

For each individual player with state s_t and realized action x_t, his state s_{t+1} in period $t+1$ will, by (2), be distributed according to $\tilde{g}_t(s_t, x_t, \tau_t|\cdot)$. Thus, it will be reasonable for the pre-action environment in period $t+1$ to follow $\sigma_{t+1} = \tau_t \odot \tilde{g}_t(\cdot, \cdot, \tau_t)$, with

$$[\tau_t \odot \tilde{g}_t(\cdot, \cdot, \tau_t)](S') = \int_{S \times X} \tau_t(ds \times dx) \cdot \tilde{g}_t(s, x, \tau_t|S'), \qquad \forall S' \in \mathcal{B}(S). \tag{6}$$

Although (6) has been intuitively reasoned from (2), we caution that logically it is part of the NG's definition rather than something derivable from the latter.

The transition from σ_t to σ_{t+1} through random action plan χ_t is best expressed by an operator. For any kernel $\chi \in \mathcal{K}(S, X)$, define operator $T_t(\chi)$ on the space $\mathcal{P}(S)$, so that

$$T_t(\chi) \circ \sigma = (\sigma \otimes \chi) \odot \tilde{g}_t(\cdot, \cdot, \sigma \otimes \chi) = \sigma \odot \chi \odot \tilde{g}_t(\cdot, \cdot, \sigma \otimes \chi), \qquad \forall \sigma \in \mathcal{P}(S). \tag{7}$$

Basically, state distribution σ and random state-dependent action plan χ first fuse to form the joint state-action distribution $\sigma \otimes \chi$ to be felt by all players. The latter's random state transitions are then guided by the kernel $\tilde{g}_t(\cdot, \cdot, \sigma_t \otimes \chi)$. Subsequently, after "averaging out" impacts of actions, the next-period state distribution will become $\sigma \odot \chi \odot \tilde{g}_t(\cdot, \cdot, \sigma \otimes \chi)$. The one-period pre-action environment transition is now representable by

$$\sigma_{t+1} = T_t(\chi_t) \circ \sigma_t = \sigma_t \odot \chi_t \odot \tilde{g}_t(\cdot, \cdot, \sigma_t \otimes \chi_t). \tag{8}$$

For periods t and t' with $t \leq t'$, as well as sequence $\chi_{[tt']} = (\chi_{t''} | t'' = t, ..., t')$ of action plans, we can iteratively define $T_{[tt']}(\chi_{[tt']})$, so that

$$T_{[tt']}(\chi_{[tt']}) \circ \sigma_t = T_{t'}(\chi_{t'}) \circ (T_{[t,t'-1]}(\chi_{[t,t'-1]}) \circ \sigma_t), \qquad \forall \sigma_t \in \mathcal{P}(S). \tag{9}$$

The left-hand side will be players' state distribution in period $t' + 1$ when they start period t with the distribution σ_t and adopt the action sequence $\chi_{[tt']}$ in the interim. Note that $T_{[tt]}(\chi_{[tt]})$ is nothing but $T_t(\chi_t)$. As a default, we let $T_{[t,t-1]}$ stand for the identity operator on $\mathcal{P}(S)$. The environment trajectory $\sigma_{[1,\bar{t}+1]}$ satisfies

$$\sigma_{[1,\bar{t}+1]} = (T_{[1,t-1]}(\chi_{[1,t-1]}) \circ \sigma_1 \mid t = 1, 2, ..., \bar{t}, \bar{t} + 1). \tag{10}$$

It is deterministic by definition.

4 The n-player Game

Let the same $\bar{t}$, S, X, $(\tilde{f}_t | t = 1, 2, ..., \bar{t})$, and $(\tilde{g}_t | t = 1, 2, ..., \bar{t})$ remain in the background. For some $n \in \mathbb{N} \setminus \{1\}$ and initial multi-state $s_1 = (s_{11}, s_{12}, ..., s_{1n}) \in S^n$, we can define an n-player game $\Gamma_n(s_1)$, in which each $s_{1m} \in S$ is player m's initial state. The game's payoffs and state evolutions are still described by the $\tilde{f}_t$'s and $\tilde{g}_t$'s, respectively. However, details are messier as outside environments vary from player to player and their evolutions are random.

For $a \in A$, where A is again a separable metric space, we use δ_a to denote the singleton Dirac measure with $\delta_a(\{a\}) = 1$. For $a = (a_1, ..., a_n) \in A^n$ where $n \in \mathbb{N}$, we use ε_a for $\sum_{m=1}^{n} \delta_{a_m}/n$, the empirical distribution generated by the vector a. We also use $\mathcal{P}_n(A)$ to

denote the space of probability measures of the type ε_a for $a \in A^n$, i.e., the space of empirical distributions generated from n samples. Now back at the game $\Gamma_n(s_1)$, suppose in period $t = 1, 2, ..., \bar{t}$, each player $m = 1, 2, ..., n$ is in state s_{tm} and takes action x_{tm}. Then, the in-action environment experienced by player 1 will be $\varepsilon_{s_{t,-1}x_{t,-1}} = \varepsilon_{((s_{t2},x_{t2}),...,(s_{tn},x_{tn}))}$. Thus, this player will receive payoff $\tilde{f}_t(s_{t1}, x_{t1}, \varepsilon_{s_{t,-1}x_{t,-1}})$ in the period, and his period-$(t+1)$ state $s_{t+1,1}$ will be sampled from the distribution $\tilde{g}_t(s_{t1}, x_{t1}, \varepsilon_{s_{t,-1}x_{t,-1}}|\cdot)$.

Suppose $\chi_{[1\bar{t}]} = (\chi_t \mid t = 1, ..., \bar{t}) \in (\mathcal{K}(S, X))^{\bar{t}}$ again describes the policy adopted by all n players. Unlike in an NG, this time $\chi_{[1\bar{t}]}$ will help to generate a stochastic as opposed to deterministic environment trajectory. To describe each one-period transition in this complex process, we rely on the kernel ${\chi_t}^n \odot {\tilde{g}_t}^n \in \mathcal{K}(S^n, S^n)$ defined by

$$({\chi_t}^n \odot {\tilde{g}_t}^n)(s|S') = \int_{X^n} {\chi_t}^n(s|dx) \cdot {\tilde{g}_t}^n(s, x|S'), \qquad \forall s \in S^n, S' \in \mathcal{B}(S^n), \tag{11}$$

where ${\chi_t}^n$ is a member of $\mathcal{K}(S^n, X^n)$ that satisfies

$${\chi_t}^n(s|X'_1 \times \cdots \times X'_n) = \Pi_{m=1}^n \chi_t(s_m|X'_m), \qquad \forall s \in S^n, X'_1, ..., X'_n \in \mathcal{B}(X), \tag{12}$$

and ${\tilde{g}_t}^n$ is a member of $\mathcal{K}(S^n \times X^n, S^n)$ that satisfies

$$\begin{gathered} {\tilde{g}_t}^n(s, x|S'_1 \times \cdots \times S'_n) = \Pi_{l=1}^n \tilde{g}_t(s_l, x_l, \varepsilon_{s_{-l}x_{-l}}|S'_l), \\ \forall (s, x) \in S^n \times X^n, S'_1, ..., S'_n \in \mathcal{B}(S). \end{gathered} \tag{13}$$

In combination, (11) can be spelled out as

$$({\chi_t}^n \odot {\tilde{g}_t}^n)(s|S'_1 \times \cdots \times S'_n) = \int_{X^n} \Pi_{m=1}^n \chi_t(s_m|dx_m) \cdot \Pi_{l=1}^n \tilde{g}_t(s_l, x_l, \varepsilon_{s_{-l}x_{-l}}|S'_l). \tag{14}$$

The above reflects that, each player m samples his action x_m from the distribution $\chi_t(s_m|\cdot)$; once all players' actions $x = (x_1, ..., x_n)$ have been determined, each player l will face his unique in-action environment $\varepsilon_{s_{-l}x_{-l}}$; thus, this player's period-$(t+1)$ state will be sampled from the distribution $\tilde{g}_t(s_l, x_l, \varepsilon_{s_{-l}x_{-l}}|\cdot)$.

When the n players start period t with a random multi-state with distribution $\pi_{nt} \in \mathcal{P}(S^n)$ and they act according to random rule $\chi_t \in \mathcal{K}(S, X)$ in the period, they will generate the joint distribution $\mu_{nt} \in \mathcal{P}(S^n \times X^n)$ of period-t multi-state and -action satisfying

$$\mu_{nt} = \pi_{nt} \otimes {\chi_t}^n. \tag{15}$$

According to (3) and (12), the above means that, for any $S' \in \mathcal{B}(S^n)$ and $X'_1, ..., X'_n \in \mathcal{B}(X)$,

$$\mu_{nt}(S' \times X'_1 \times \cdots \times X'_n) = \int_{S'} \pi_{nt}(ds) \cdot {\chi_t}^n(s|X'_1 \times \cdots \times X'_n) = \int_{S'} \pi_{nt}(ds) \cdot \Pi_{m=1}^n \chi_t(s_m|X'_m). \tag{16}$$

Clearly, (15) corresponds to (5) in the NG situation.

By (11), the period-$(t+1)$ multi-state distribution $\mu_{nt} \odot \tilde{g}_t^{\,n} \in \mathcal{P}(S^n)$ will follow

$$(\mu_{nt} \odot \tilde{g}_t^{\,n})(S') = \int_{S^n \times X^n} \mu_{nt}(ds \times dx) \cdot \tilde{g}_t^{\,n}(s, x|S'), \qquad \forall S' \in \mathcal{B}(S^n). \tag{17}$$

Combining (15) and (17), we can see that the one-period transition between multi-states is

$$\pi_{n,t+1} = (\pi_{nt} \otimes \chi_t^{\,n}) \odot \tilde{g}_t^{\,n} = \pi_{nt} \odot \chi_t^{\,n} \odot \tilde{g}_t^{\,n}. \tag{18}$$

Note (18) is the n-player game's answer to the NG's (8). Similar to (9), for $t \leq t'$, the distribution $\pi_{nt'}$ of period-t' multi-state $s_{t'}$ is given by

$$\pi_{nt'} = \pi_{nt} \odot \Pi_{t''=t}^{t'-1}(\chi_{t''}^{\,n} \odot \tilde{g}_{t''}^{\,n}). \tag{19}$$

When the initial multi-state s_1 is randomly drawn from distribution π_{n1}, the entire trajectory $\pi_{n,[1,\bar{t}+1]} = (\pi_{nt}|t = 1, 2, ..., \bar{t}, \bar{t}+1)$ of the n-player game's multi-state distributions can be written as

$$\pi_{n,[1,\bar{t}+1]} = (\pi_{n1} \odot \Pi_{t'=1}^{t-1}(\chi_{t'}^{\,n} \odot \tilde{g}_{t'}^{\,n})|t = 1, 2, .., \bar{t}, \bar{t}+1). \tag{20}$$

When all players' states are sampled from some $\sigma_1 \in \mathcal{P}(S)$, we still have (20) as the trajectory for multi-state distributions, but with $\pi_{n1} = \sigma_1^{\,n}$. When recognizing $\pi_{n1} = \delta_{s_1}$, the Dirac measure in $\mathcal{P}(S^n)$ that assigns the full weight to s_1, (20) will help describe the evolution of the multi-state distribution for the n-player game $\Gamma_n(s_1)$, much like (10) did for $\Gamma(\sigma_1)$.

5 Convergence of Aggregate Environments

Even before touching upon notions like cumulative payoffs and equilibria, we can already introduce an interesting link between finite games and NGs. It is in terms of an asymptotic relationship between a sequence $\pi_{n,[t,\bar{t}+1]} = (\pi_{nt'}|t' = t, t+1, ..., \bar{t}+1)$ of multi-state distributions in n-player games and a sequence $\sigma_{[t,\bar{t}+1]} = (\sigma_{t'}|t' = t, t+1, ..., \bar{t}+1)$ of state distributions in their NG counterparts. The message is that, when starting from similar environments in period t and adopting the same action plan from that period on, stochastic environment paths experienced by large finite games will not drift too much away from the NG's deterministic environment trajectory. We refrain from using the word convergence because the $\pi_{nt'}$'s reside in different spaces for different n's.

First, we propose the concept asymptotic resemblance in order to precisely describe the way in which members in a sequence of probability measures increasingly resemble the

products of a given measure. For a separable metric space A, the space $\mathcal{P}(A)$ is metrized by the Prohorov metric ρ_A, which induces the weak topology on it. At fixed $n \in \mathbb{N}$, the map $\varepsilon_{(\cdot)}$ from A^n to $\mathcal{P}_n(A) \subseteq \mathcal{P}(A)$ is continuous. Therefore, for any $p \in \mathcal{P}(A)$ and $\epsilon > 0$, the set $\{a \in A^n | \rho_A(\varepsilon_a, p) < \epsilon\}$ is an open subset of A^n and thus a member of $\mathcal{B}(A^n)$.

Definition 1 *For a separable metric space A, suppose $p \in \mathcal{P}(A)$ and for each $n \in \mathbb{N}$, $q_n \in \mathcal{P}(A^n)$. We say that sequence q_n asymptotically resembles the sequence p^n made up of p's n-th order products $p \times \cdots \times p$, if for any $\epsilon > 0$ and n that is large enough,*

$$q_n(\{a \in A^n | \rho_A(\varepsilon_a, p) < \epsilon\}) > 1 - \epsilon.$$

Definition 1 says that sequence q_n will asymptotically resemble the sequence p^n of product measures when the empirical distribution ε_a of a random vector $a = (a_1, ..., a_n)$, sampled from q_n, is highly likely to be close to p as n approaches $+\infty$. This resemblance notion is consistent with Prohorov's theorem (Parthasarathy [26], Theorem II.7.1), whose weak version is presented as Lemma 2 in Appendix A. Due to it, any sequence $(p')^n$ will asymptotically resemble the sequence p^n if and only if $p' = p$.

Some results related to the resemblance concept have been placed in Appendix A. Lemma 3 stems from Dvoretzky, Kiefer, and Wolfolwitz's [10] inequality and makes the convergence in Lemma 2 uniform in the chosen probability p. According to Lemma 4, the tampering of one component within any n-long vector $a \in A^n$ would not much alter ε_a. It is therefore natural for Lemma 5 to state that the resemblance of q_n to p^n would lead to that of the A^{n-1}-marginal $q_n|_{A^{n-1}}$ to p^{n-1}. Lemma 6 says that the above would also lead to the asymptotic resemblance of $p' \times q_{n-1}$ to p^n for any p'. So in general there can be nothing substantial regarding the relationship between the A-marginals $q_n|_A$ and p. Finally, Lemma 7 shows that asymptotic resemblance is preserved under the projection of $A \times B$ into A.

The following one-step result states that asymptotic resemblance concerning pre-action environments is translatable into that concerning in-action environments; also, the same resemblance is preserved after undergoing one single step in a game.

Proposition 1 *Let state distribution $\sigma \in \mathcal{P}(S)$, random state-dependent action plan $\chi \in \mathcal{K}(S, X)$, and state-transition kernel $g \in \mathcal{G}(S, X)$, with the latter enjoying the continuity of $g(s, x, \tau)$ in the joint state-action distribution τ at an (s, x)-independent rate. Also, multi-state distribution $\pi_n \in \mathcal{P}(S^n)$ for each $n \in \mathbb{N}$. Suppose further that the sequence π_n asymptotically resembles the sequence σ^n. Then,*

(i) the sequence $\pi_n \otimes \chi^n$ will asymptotically resemble the sequence $(\sigma \otimes \chi)^n$, and

(ii) the sequence $\pi_n \odot \chi^n \odot g^n$ will asymptotically resemble the sequence $(\sigma \odot \chi \odot g(\cdot, \cdot, \sigma \otimes \chi))^n$.

Indeed, (ii) remains valid under mild contamination. That is, for any $(s, x) \in S \times X$,

(iii) the sequence $(\delta_{sx} \times (\pi_{n-1} \otimes \chi^{n-1})) \odot g^n$ will asymptotically resemble the sequence $(\sigma \odot \chi \odot g(\cdot, \cdot, \sigma \otimes \chi))^n$ at a rate independent of the chosen (s, x).

Proposition 1 is one of our two most technical results. Its proof invokes both Prohorov's theorem (Parthasarathy [26], Theorem II.7.1) on the convergence of empirical distributions and for parts (ii) and (iii), Dvoretzky, Kiefer, and Wolfolwitz's [10] inequality which provides the uniformity of such convergence. In the proposition, part (i) stresses the passibility from convergence of pre-action environments to that of same-period in-action environments, see (5) and (15); part (ii) further points out that convergence in next-period pre-action environments will follow suit, see (8) and (18); also, part (iii) will be useful when we take the view point from one single player..

To take advantage of Proposition 1, we now assume the equi-continuity of the state transitions with respect to in-action environments.

Assumption 1 *Each transition kernel $\tilde{g}_t(s, x, \tau)$ is continuous in τ at an (s, x)-independent rate. That is, for any in-action environment $\tau \in \mathcal{P}(S \times X)$ and $\epsilon > 0$, there is $\delta > 0$, such that for any $\tau' \in \mathcal{P}(S \times X)$ satisfying $\rho_{S \times X}(\tau, \tau') < \delta$ and any $(s, x) \in S \times X$,*

$$\rho_S(\tilde{g}_t(s, x, \tau), \tilde{g}_t(s, x, \tau')) < \epsilon.$$

We are in a position to derive this section's main result. It states that, when an NG and its finite counterparts evolve under the same action plan, environment pathways of large finite games, though stochastic, will resemble the deterministic pathway of the NG.

Theorem 1 *Let a policy profile $\chi_{[t\bar{t}]} \in (\mathcal{K}(S, X))^{\bar{t}-t+1}$ for periods $t, t+1, ..., \bar{t}$ be given. When $s_t = (s_{t1}, ..., s_{tn})$ has a distribution π_{nt} that asymptotically resembles ${\sigma_t}^n$, the series $(\pi_{nt} \odot \Pi_{t''=t}^{t'-1}(\chi_{t''}^n \odot \tilde{g}_{t''}^n) \mid t' = t, t+1, ..., \bar{t}, \bar{t}+1)$ will asymptotically resemble $((T_{[t,t'-1]}(\chi_{[t,t'-1]}) \circ \sigma_t)^n \mid t' = t, t+1, ..., \bar{t}, \bar{t}+1)$ as well. That is, for any $\epsilon > 0$ and any n that is large enough,*

$$[\pi_{nt} \odot \Pi_{t''=t}^{t'-1}(\chi_{t''}^n \odot \tilde{g}_{t''}^n)](\tilde{A}_{nt'}(\epsilon)) > 1 - \epsilon, \qquad \forall t' = t, t+1, ..., \bar{t}+1,$$

where for each t', the set of multi-states $\tilde{A}_{nt'}(\epsilon) \in \mathcal{B}(S^n)$ is such that,

$$\rho_S(\varepsilon_{s_{t'}}, T_{[t,t'-1]}(\chi_{[t,t'-1]}) \circ \sigma_t) < \epsilon, \qquad \forall s_{t'} \in \tilde{A}_{nt'}(\epsilon).$$

Suppose an NG starts period t with pre-action environment σ_t and a slew of finite games start the period with pre-action environments that are ever nearly sampled from σ_t. Let the evolution of both types of games be guided by players acting according to the same policy profile $\chi_{[t\bar{t}]}$. Then, as the numbers of players n involved in finite games grow indefinitely, Theorem 1 predicts for ever less chances for the finite games' period-t' environments $\varepsilon_{s_{t'}}$ to be even slightly away from the NG's deterministic period-t' environment $T_{[t,t'-1]}(\chi_{[t,t'-1]})\circ\sigma_t$. For some fixed $\sigma_1 \in \mathcal{P}(S)$, we can plug $t = 1$ and $\pi_{n1} = {\sigma_1}^n$ into Theorem 1. Then, we will obtain the proximity between $\sigma^n_{[1,\bar{t}+1]} = ({\sigma_t}^n | t = 1, 2, ..., \bar{t}, \bar{t}+1)$ and $\pi_{n,[1,\bar{t}+1]} = (\pi_{nt} | t = 1, 2, ..., \bar{t}, \bar{t}+1)$ for large n's, where every $\sigma_t = T_{[1,t-1]}(\chi_{[1,t-1]}) \circ \sigma_1$ and every $\pi_{nt} = {\sigma_1}^n \odot \Pi_{t'=1}^{t-1}(\chi_{t'}{}^n \odot \tilde{g}_{t'}{}^n)$. In view of (10) and (20), this means that when large games sample their initial states from an NG's starting distribution σ_1, the former games' state-distribution trajectories will remain close to that of the latter game.

Our confinement so far to discrete spaces S and X arises mainly from the need to deal with non-product joint probabilities of the form $p \otimes \kappa$; see (3). In Yang [34], where random state transitions and random action plans were modeled through independently generated shocks, only results pertaining to product-form probabilities $p \times q$, where q is an ordinary rather than conditional probability, were needed. Because of this, known properties like Propositions III.4.4 and III.4.6 of Ethier and Kurtz [11] could be put to good use. Results there could thus be based on complete state and shock spaces. In contrast, if we were to consider more general spaces here, we would face the presently unsurmountable challenge of passing the closeness between measures p and p_i for $i = 1, 2, ..., n$ onto that between p^n and $\prod_{i=1}^n p_i$ when n itself tends to infinity.

6 NG and Finite-game Equilibria

We present this paper's main result that an NG equilibrium, though oblivious of past history and blind to other players' states, will generate minimal regrets when adopted by players in large finite games. First, we introduce equilibrium concepts used in both types of games.

6.1 Equilibria in NG

In defining the NG $\Gamma(\sigma_1)$'s equilibria, we subject a candidate policy profile to one-time deviation of a single player, who is by default infinitesimal in influence. Note the deviation will not alter the environment trajectory corresponding to the candidate profile. With this

understanding, we define $v_t(s_t, \xi_{[t\bar{t}]}, \sigma_t, \chi_{[t\bar{t}]})$ as the total expected payoff a player can receive from period t to $\bar{t}$, when he starts with state $s_t \in S$ and adopts action plan $\xi_{[t\bar{t}]} = (\xi_t, ..., \xi_{\bar{t}}) \in (\mathcal{K}(S,X))^{\bar{t}-t+1}$ throughout, while other players form initial pre-action environment $\sigma_t \in \mathcal{P}(S)$ and adopt policy profile $\chi_{[t\bar{t}]} = (\chi_t, ..., \chi_{\bar{t}}) \in (\mathcal{K}(S,X))^{\bar{t}-t+1}$ throughout. As a terminal condition, we certainly have

$$v_{\bar{t}+1}(s_{\bar{t}+1}, \sigma_{\bar{t}+1}) = 0. \tag{21}$$

For $t = \bar{t}, \bar{t}-1, ..., 1$, we have the recursive relationship

$$\begin{aligned} v_t(s_t, \xi_{[t\bar{t}]}, \sigma_t, \chi_{[t\bar{t}]}) = \textstyle\int_X \xi_t(s_t|dx_t) \cdot [\tilde{f}_t(s_t, x_t, \sigma_t \otimes \chi_t) \\ + \textstyle\int_S \tilde{g}_t(s_t, x_t, \sigma_t \otimes \chi_t|ds_{t+1}) \cdot v_{t+1}(s_{t+1}, \xi_{[t+1,\bar{t}]}, T_t(\chi_t) \circ \sigma_t, \chi_{[t+1,\bar{t}]})]. \end{aligned} \tag{22}$$

This is because the player's action is guided in a random fashion by ξ_t, its payoff is determined by $\tilde{f}_t$, its state evolution is governed by $\tilde{g}_t$, and its future payoff is supplied by v_{t+1}; also, after undergoing the commonly adopted action plan χ_t, the period-$(t+1)$ pre-action environment σ_{t+1} will be $T_t(\chi_t) \circ \sigma_t$ as shown in (8). The choice of ξ_t affects the current player's period-t action x_t, his period-$(t+1)$ state s_{t+1}, and his future state-action trajectory. However, the change at this negligible player does not alter the period-t in-action environment $\sigma_t \otimes \chi_t$ as listed in (5) or any environment in the future. This is the main reason why NGs are easier to handle than their finite-player counterparts.

Now, we deem policy $\chi_{[1\bar{t}]} \in (\mathcal{K}(S,X))^{\bar{t}}$ a Markov equilibrium for the game $\Gamma(\sigma_1)$ when, for every $t = 1, 2, ..., \bar{t}$ and $\xi_t \in \mathcal{K}(S,X)$,

$$v_t(s_t, \chi_{[t\bar{t}]}, \sigma_t, \chi_{[t\bar{t}]}) \geq v_t(s_t, (\xi_t, \chi_{[t+1,\bar{t}]}), \sigma_t, \chi_{[t\bar{t}]}), \qquad \forall s_t \in S, \tag{23}$$

where

$$\sigma_t = T_{[1,t-1]}(\chi_{[1,t-1]}) \circ \sigma_1. \tag{24}$$

That is, policy $\chi_{[1\bar{t}]}$ will be regarded an equilibrium when no player can be better off by unilaterally deviating to any alternative plan $\xi_t \in \mathcal{K}(S,X)$ in any single period t. The definition of σ_t in (24) underscores the evolution of the deterministic environment trajectory following the adoption of action plan $\chi_{[1\bar{t}]}$ by almost all players.

6.2 ϵ-Equilibria in n-player Games

For an n-player game, let $v_{nt}(s_{t1}, \xi_{[t\bar{t}]}, \varepsilon_{s_{t,-1}}, \chi_{[t\bar{t}]})$ be the total expected payoff player 1 can receive from period t to $\bar{t}$, when he starts with state $s_{t1} \in S$ and adopts action plan

$\xi_{[t\bar{t}]} \in (\mathcal{K}(S,X))^{\bar{t}-t+1}$ throughout, while other players form initial empirical state distribution $\varepsilon_{s_{t,-1}} = \varepsilon_{(s_{t2},...,s_{tn})} \in \mathcal{P}_{n-1}(S)$ and adopt action plan $\chi_{[t\bar{t}]} \in (\mathcal{K}(S,X))^{\bar{t}-t+1}$ throughout. As a terminal condition, we have

$$v_{n,\bar{t}+1}(s_{\bar{t}+1,1}, \varepsilon_{s_{\bar{t}+1,-1}}) = 0. \tag{25}$$

For $t = \bar{t}, \bar{t}-1, ..., 1$, we have the recursive relationship

$$\begin{array}{l} v_{nt}(s_{t1}, \xi_{[t\bar{t}]}, \varepsilon_{s_{t,-1}}, \chi_{[t\bar{t}]}) = \int_X \xi_t(s_{t1}|dx_{t1}) \cdot \int_{X^{n-1}} \chi_t^{n-1}(s_{t,-1}|dx_{t,-1}) \times \\ \quad \times [\tilde{f}_t(s_{t1}, x_{t1}, \varepsilon_{s_{t,-1}x_{t,-1}}) + \int_{S^n} \tilde{g}_t^n(s_t, x_t|ds_{t+1}) \cdot v_{n,t+1}(s_{t+1,1}, \xi_{[t+1,\bar{t}]}, \varepsilon_{s_{t+1},-1}, \chi_{[t+1,\bar{t}]})], \end{array} \tag{26}$$

where the meaning of $\chi_t^{n-1}(s_{t,-1}|dx_{t,-1})$ follows from (12) and that of $\tilde{g}_t^n(s_t, x_t|ds_{t+1})$ follows from (13). Note (26) differs substantially from its NG counterpart (22). With only a finite number of players, player 1's one-time choice ξ_t not only affects his own future actions and states as before, but differently, starting from the altered in-action environment $\varepsilon_{s_t x_t}$, it also impacts the entire future trajectory of all other players. Note $\varepsilon_{s_t x_t}$ impacts the generation of $s_{t+1} = (s_{t+1,1}, ..., s_{t+1,n})$ in its projections to n different $(n-1)$-dimensional spaces, as according to (13), $\int_{S^n} \tilde{g}_t^n(s_t, x_t|ds_{t+1})$ amounts to $\Pi_{m=1}^n \int_S \tilde{g}_t(s_{tm}, x_{tm}, \varepsilon_{s_{t,-m}x_{t,-m}}|ds_{t+1,m})$.

For each $n \in \mathbb{N} \setminus \{1\}$, let $\hat{\pi}_{n-1,[1\bar{t}]} = (\hat{\pi}_{n-1,t} \mid t = 1, ..., \bar{t}) \in (\mathcal{P}(S^{n-1}))^{\bar{t}}$ be a series of other-player multi-state distributions. For $\epsilon \geq 0$, we deem $\chi_{[1\bar{t}]} = (\chi_t \mid t = 1, ..., \bar{t}) \in (\mathcal{K}(S,X))^{\bar{t}}$ an ϵ-Markov equilibrium for the game family $(\Gamma_n(s_1) \mid s_1 \in S^n)$ in the sense of $\hat{\pi}_{n-1,[1\bar{t}]}$ when, for every $t = 1, ..., \bar{t}$, $\xi_{[t\bar{t}]} \in (\mathcal{K}(S,X))^{\bar{t}-t+1}$, and $s_{t1} \in S$,

$$\begin{array}{l} \int_{S^{n-1}} \hat{\pi}_{n-1,t}(ds_{t,-1}) \cdot v_{nt}(s_{t1}, \chi_{[t\bar{t}]}, \varepsilon_{s_{t,-1}}, \chi_{[t\bar{t}]}) \\ \qquad \geq \int_{S^{n-1}} \hat{\pi}_{n-1,t}(ds_{t,-1}) \cdot v_{nt}(s_{t1}, \xi_{[t\bar{t}]}, \varepsilon_{s_{t,-1}}, \chi_{[t\bar{t}]}) - \epsilon. \end{array} \tag{27}$$

That is, action plan $\chi_{[1\bar{t}]}$ will be an ϵ-Markov equilibrium in the sense of $\hat{\pi}_{n-1,[1\bar{t}]}$ when under the plan's guidance, the average payoff from any period t and player-1 state s_{t1} on cannot be improved by more than ϵ through any unilateral deviation, where the "average" is based on other players' multi-state $s_{t,-1}$ being sampled from the distribution $\hat{\pi}_{n-1,t}$. Note (27) differs from (23) also in that its unilateral deviation need not be one-time.

6.3 Main Transient Result

Before moving on, we need the single-period payoff functions $\tilde{f}_t$ to be continuous.

Assumption 2 *Each payoff function $\tilde{f}_t(s,x,\tau)$ is continuous in the in-action environment τ at an (s,x)-independent rate. That is, for any $\tau \in \mathcal{P}(S \times X)$ and $\epsilon > 0$, there is $\delta > 0$,*

such that for any $\tau' \in \mathcal{P}(S \times X)$ *satisfying* $\rho_{S\times X}(\tau, \tau') < \delta$ *and any* $(s, x) \in S \times X$,

$$\mid \tilde{f}_t(s, x, \tau) - \tilde{f}_t(s, x, \tau') \mid < \epsilon.$$

Now we show the convergence of finite-game value functions to their NG counterpart, the proof of which is quite technical as well, and calls upon parts (i) and (iii) of Proposition 1.

Proposition 2 *For any* $t = 1, 2, ..., \bar{t} + 1$, *let* $\sigma_t \in \mathcal{P}(S)$ *and* $\hat{\pi}_{n-1,t} \in \mathcal{P}(S^{n-1})$ *for each* $n \in \mathbb{N}$. *Suppose the sequence* $\hat{\pi}_{n-1,t}$ *asymptotically resembles the sequence* ${\sigma_t}^{n-1}$. *Then for any* $\chi_{[t\bar{t}]} \in (\mathcal{K}(S, X))^{\bar{t}-t+1}$, *the sequence* $\int_{S^{n-1}} \hat{\pi}_{n-1,t}(ds_{t,-1}) \cdot v_{nt}(s_{t1}, \xi_{[t\bar{t}]}, \varepsilon_{s_{t,-1}}, \chi_{[t\bar{t}]})$ *will converge to* $v_t(s_{t1}, \xi_{[t\bar{t}]}, \sigma_t, \chi_{[t\bar{t}]})$ *at a rate that is independent of both* $s_{t1} \in S$ *and* $\xi_{[t\bar{t}]} \in (\mathcal{K}(S, X))^{\bar{t}-t+1}$.

Combining (23) and (27), as well as Proposition 2, we can come to the main result.

Theorem 2 *For some* $\sigma_1 \in \mathcal{P}(S)$, *suppose* $\chi_{[1\bar{t}]} = (\chi_t \mid t = 1, 2, ..., \bar{t}) \in (\mathcal{K}(S, X))^{\bar{t}}$ *is a Markov equilibrium of NG* $\Gamma(\sigma_1)$. *Also, suppose* $\hat{\pi}_{n-1,[1\bar{t}]} = (\hat{\pi}_{n-1,t} | t = 1, 2, ..., \bar{t}) \in (\mathcal{P}(S^{n-1}))^{\bar{t}}$ *is such that the sequence* $\hat{\pi}_{n-1,t}$ *asymptotically resembles the sequence* ${\sigma_t}^{n-1}$ *for each* t, *where* $\sigma_t = T_{[1,t-1]}(\chi_{[1,t-1]}) \circ \sigma_1$. *Then, for* $\epsilon > 0$ *and large enough* $n \in \mathbb{N}$, *the given* $\chi_{[1\bar{t}]}$ *is also an* ϵ*-Markov equilibrium for the game family* $(\Gamma_n(s_1) \mid s_1 \in S^n)$ *in the sense of* $\hat{\pi}_{n-1,[1\bar{t}]}$.

The theorem says that players in a large finite game can agree on an NG equilibrium and expect to lose little on average, as long as the other-player multi-state distribution $\hat{\pi}_{n-1,t}$ on which "average" is based is similar to the product form ${\sigma_t}^{n-1}$, where $\sigma_t = T_{[1,t-1]}(\chi_{[1,t-1]}) \circ \sigma_1$ is the corresponding NG's predictable equilibrium state distribution for the same period. As to whether reasonable $\hat{\pi}_{n-1,[1\bar{t}]} = (\hat{\pi}_{n-1,t} | t = 1, 2, ..., \bar{t})$ exists to satisfy this condition, the answer is affirmative. The next section is dedicated to this point.

7 The Condition in Theorem 2

We now present examples where the key condition in Theorem 2 can be true. In all of them, we let the initial other-player multi-state distribution $\hat{\pi}_{n-1,1} = {\sigma_1}^{n-1} = {\sigma_1}^n|_{S^{n-1}}$. That is, we let players' initial states in n-player games be randomly drawn from the NG's initial state distribution σ_1. Now we discuss what can happen in periods $t = 2, 3, ..., \bar{t}$.

7.1 Two Possibilities

First, we can let each $\hat{\pi}_{n-1,t} = {\sigma_t}^{n-1}$. It has been discussed right after Definition 1 that the sequence ${\sigma_t}^{n-1}$ asymptotically resembles itself. So this choice satisfies the condition in Theorem 2. This would correspond to the case where players in large finite games take the "lazy" approach of using independent draws on the NG state distribution to assess their opponents' states. Note this is reasonable due to the common initial condition for both types of games and Theorem 1.

Second, we can let each $\hat{\pi}_{n-1,t} = \pi_{nt}|_{S^{n-1}}$, where

$$\pi_{nt} = {\sigma_1}^n \odot \Pi_{t'=1}^{t-1}(\chi_{t'}{}^n \odot \tilde{g}_{t'}{}^n). \tag{28}$$

According to (19), π_{nt} stands for players' multi-state distribution in period t in an n-player game when their initial states are randomly drawn from the distribution σ_1 and then from period 1 onward players all follow through with the NG equilibrium $\chi_{[1\bar{t}]}$. Since the sequence ${\sigma_1}^n$ asymptotically resembles itself, Theorem 1 will ascertain the asymptotic resemblance of π_{nt} to ${\sigma_t}^n$. Then, Lemma 5 in Appendix A will lead to the asymptotic resemblance of $\hat{\pi}_{n-1,t}$ to ${\sigma_t}^{n-1}$. So this choice would satisfy Theorem 2's condition as well. Also, its meaning is clear—here players in large finite games use precise assessments on what other players' states might be had they followed the NG equilibrium all along.

7.2 Refinement and a Third Choice

Note that $\hat{\pi}_{n-1,t}$ has not countenanced the possibility in which a player involves his own state s_{t1} in the estimation of the other-player multi-state $s_{t,-1}$. We now show that this is possible at least when the state space S is finite. In that case, we can upgrade the $\hat{\pi}_{n-1,t} \in \mathcal{P}(S^{n-1})$ in Proposition 2 to $\hat{\pi}_{n-1,t}(\cdot) = (\hat{\pi}_{n-1,t}(s_{t1}|\cdot)|s_{t1} \in S) \in (\mathcal{P}(S^{n-1}))^S$ and obtain the convergence of $\int_{S^{n-1}} \hat{\pi}_{n-1,t}(s_{t1}|ds_{t,-1}) \cdot v_{nt}(s_{t1}, \xi_{[t\bar{t}]}, \varepsilon_{s_{t,-1}}, \chi_{[t\bar{t}]})$ to $v_t(s_{t1}, \xi_{[t\bar{t}]}, \sigma_t, \chi_{[t\bar{t}]})$ at an s_{t1}-independent rate. This will lead us to the following extended version of Theorem 2.

Theorem 3 *Suppose $\sigma_{[1,\bar{t}+1]}$ and $\chi_{[1\bar{t}]}$ are all the same as in Theorem 2. Also, suppose $\hat{\pi}_{n-1,[1\bar{t}]}(\cdot) = (\hat{\pi}_{n-1,t}(s_{t1}|\cdot)|t = 1, 2, ..., \bar{t},\ s_{t1} \in S) \in ((\mathcal{P}(S^{n-1}))^S)^{\bar{t}}$ is such that the sequence $\hat{\pi}_{n-1,t}(s_{t1}|\cdot)$ asymptotically resembles the sequence ${\sigma_t}^{n-1}$ for each t and s_{t1}. Then, for $\epsilon > 0$ and large enough $n \in \mathbb{N}$, for every $t = 1, ..., \bar{t}$, $\xi_{[t\bar{t}]} \in (\mathcal{K}(S, X))^{\bar{t}-t+1}$, and $s_{t1} \in S$,*

$$\begin{aligned} &\int_{S^{n-1}} \hat{\pi}_{n-1,t}(s_{t1}|ds_{t,-1}) \cdot v_{nt}(s_{t1}, \chi_{[t\bar{t}]}, \varepsilon_{s_{t,-1}}, \chi_{[t\bar{t}]}) \\ &\qquad \geq \int_{S^{n-1}} \hat{\pi}_{n-1,t}(s_{t1}|ds_{t,-1}) \cdot v_{nt}(s_{t1}, \xi_{[t\bar{t}]}, \varepsilon_{s_{t,-1}}, \chi_{[t\bar{t}]}) - \epsilon. \end{aligned}$$

For it to satisfy the condition in Theorem 3, we can still let $\hat{\pi}_{n-1,[1\bar{t}]}(\cdot)$ be the same as in the aforementioned two examples, in which the newly added s_{t1}-dependence is mute. But a third choice would allow each player a full-fledged Bayesian update on other players' states.

In this third choice, we still use (28) to define π_{nt}. Then, as long as $\sigma_t(s_{t1}) > 0$, we let

$$\hat{\pi}_{n-1,t}(s_{t1}|\cdot) = \pi_{nt,S}|_{S^{n-1}}(s_{t1}|\cdot), \tag{29}$$

the other-player multi-state distribution derivable from π_{nt} when conditioned on the current player's state s_{t1}; otherwise, we simply let $\hat{\pi}_{n-1,t} = \pi_{nt}|_{S^{n-1}}$ just as in the second example. Note the marginal $\pi_{nt}|_S$ is defined by

$$\pi_{nt}|_S(\{s_{t1}\}) = \pi_{nt}(\{s_{t1}\} \times S^{n-1}), \qquad \forall s_{t1} \in S, \tag{30}$$

and each conditional distribution $\pi_{nt,S}|_{S^{n-1}}(s_{t1}|\cdot)$ is defined by

$$\pi_{nt,S}|_{S^{n-1}}(s_{t1}|S') = \frac{\pi_{nt}(\{s_{t1}\} \times S')}{\pi_{nt}|_S(\{s_{t1}\})} = \frac{\pi_{nt}(\{s_{t1}\} \times S')}{\pi_{nt}(\{s_{t1}\} \times S^{n-1})}, \qquad \forall S' \in \mathcal{B}(S^{n-1}), \tag{31}$$

when the denominator is strictly positive and an arbitrary value otherwise.

7.3 Symmetry Makes it Work

The lone fact that π_{nt} asymptotically resembles ${\sigma_t}^n$ is actually quite far from being able to dictate the asymptotic resemblance of the thus defined $\hat{\pi}_{n-1,t}(s_{t1}|\cdot)$ to ${\sigma_t}^{n-1}$. Note that for a general q_n resembling some p^n, Lemma 6 in Appendix A has all but ruled out the convergence of $\pi_n|_A$ to p, let alone the asymptotic resemblance of $q_{n,A}|_{A^{n-1}}$ to p^{n-1}. Fortunately, π_{nt} still enjoys the additional feature of being symmetric.

For any $n \in \mathbb{N}$, let Ψ_n be the set of all n-dimensional permutations. That is, each $\psi \in \Psi_n$ makes $(\psi(1), ..., \psi(n))$ a permutation of $(1, ..., n)$. For a given $\psi \in \Psi_n$, let us suppose $\psi a = (a_{\psi(1)}, ..., a_{\psi(n)})$ for any $a = (a_1, ..., a_n) \in A^n$, and then $\psi A' = \{\psi a | a \in A'\}$ for any $A' \subseteq A^n$. Note that, due to its innately symmetric definition, $\mathcal{B}(A^n)$ is automatically symmetric in the sense that $\mathcal{B}(A^n) = \{\psi A' | A' \in \mathcal{B}(A^n)\}$ for any $\psi \in \Psi_n$.

Definition 2 *For $n \in \mathbb{N}$ and separable metric space A, we say $q_n \in \mathcal{P}(A^n)$ symmetric if*

$$q_n(A') = q_n(\psi A'), \qquad \forall \psi \in \Psi_n,\ A' \in \mathcal{B}(A^n).$$

We have the much needed result that asymptotic resemblance of q_n to p^n does lead to the convergence of $q_n|_A$ to p when q_n is symmetric. This is in stark contrast with Lemma 6.

Proposition 3 *Let A be a discrete metric space and $q_n \in \mathcal{P}(A^n)$ for every $n \in \mathbb{N}$ be symmetric. Suppose the sequence q_n asymptotically resembles the sequence p^n. Then, the sequence $q_n|_A$ will converge to p, namely, $\lim_{n\to+\infty} q_n|_A(\{a\}) = p(\{a\})$ for every $a \in A$.*

This then results in the resemblance of $q_{n,A}|_{A^{n-1}}$ to p^{n-1}.

Proposition 4 *Let A be a discrete metric space and $q_n \in \mathcal{P}(A^n)$ for every $n \in \mathbb{N}$ be symmetric. Suppose the sequence q_n asymptotically resembles the sequence p^n. Then, the sequence $q_{n,A}|_{A^{n-1}}(a|\cdot)$ will asymptotically resemble the sequence p^{n-1} for any $a \in A$ with $p(\{a\}) > 0$.*

Note that π_{n1}, being equal to ${\sigma_1}^n$, is symmetric. As suggested by (28), the operation it has to go through to arrive to π_{nt} is also symmetric. Hence, π_{nt} is symmetric. Therefore, by Proposition 3, the marginal probability $\pi_{nt}|_S$ as defined in (30) would converge to the NG state distribution σ_t; thus, the conditional distribution $\pi_{nt,S}|_{S^{n-1}}(s_{t1}|\cdot)$ as defined in (31) would be well defined when $\sigma_t(s_{t1}) > 0$. Then, Proposition 4 can guarantee that $\hat{\pi}_{n-1,t}(s_{t1}|\cdot)$ as defined in (29) would asymptotically resemble ${\sigma_t}^{n-1}$ and hence help to facilitate the condition needed for Theorem 3. The above suggests that, even when players exercise the most accurate Bayesian updates on other players' states using their own state information, they will not discern much regret on average by adhering to the NG equilibrium.

8 A Stationary Situation

Now we study an infinite-horizon model with stationary features. To this end, we keep S and X, but let there be a discount factor $\bar{\alpha} \in [0, 1)$. There is a payoff function $\tilde{f}$ which meets the basic measurability and boundedness requirements, so that $\tilde{f}_t = \bar{\alpha}^{t-1} \cdot \tilde{f}$ for $t = 1, 2, ...$ Let us use $\bar{f}$ for the bound $\bar{f}_1$ that appeared in (1). In addition, there is a state transition kernel $\tilde{g} \in \mathcal{G}(S, X)$, so that $\tilde{g}_t = \tilde{g}$ for $t = 1, 2, ...$ For $\chi \in \mathcal{K}(S, X)$, denote by $T(\chi)$ the operator on $\mathcal{P}(S)$, so that for any $\sigma \in \mathcal{P}(S)$,

$$T(\chi) \circ \sigma = \sigma \odot \chi \odot \tilde{g}(\cdot, \cdot, \sigma \otimes \chi). \tag{32}$$

Thus, state transition has been made stationary by the stationarity of $\tilde{g}$.

Denote the stationary nonatomic game formed from the above S, X, $\bar{\alpha}$, $\tilde{f}$, and $\tilde{g}$ by Γ^∞. It helps to first study the corresponding games Γ^t that terminate in periods $t + 1$, for $t = 0, 1, ...$ Now let $v^t(s, \xi_{[1t]}, \sigma, \chi_{[1t]})$ be the total expected payoff a player can receive in

game Γ^t, when he starts at state $s \in S$ in period 1 and adopts action plan $\xi_{[1t]} \in (\mathcal{K}(S,X))^t$ from period 1 to t, while all other players form state distribution $\sigma \in \mathcal{P}(S)$ in the beginning and act according to $\chi_{[1t]} \in (\mathcal{K}(S,X))^t$ from period 1 to t. As a terminal condition, we have $v^0(s,\sigma) = 0$. Also, for $t = 1, 2, ...,$

$$\begin{aligned} v^t(s, \xi_{[1t]}, \sigma, \chi_{[1t]}) = \textstyle\int_X \xi_1(s|dx) \cdot [\tilde{f}(s, x, \sigma \otimes \chi_1) \\ +\bar{\alpha} \cdot \textstyle\int_S \tilde{g}(s, x, \sigma \otimes \chi_1 | ds') \cdot v^{t-1}(s', \xi_{[2t]}, T(\chi_1) \circ \sigma, \chi_{[2t]})]. \end{aligned} \tag{33}$$

Using the terminal condition and (33), we can inductively show that

$$\mid v^{t+1}(s, \xi_{[1,t+1]}, \sigma, \chi_{[1,t+1]}) - v^t(s, \xi_{[1t]}, \sigma, \chi_{[1t]}) \mid \leq \bar{\alpha}^t \cdot \bar{f}. \tag{34}$$

Given $s \in S$, $\xi_{[1\infty]} = (\xi_1, \xi_2, ...) \in (\mathcal{K}(S,X))^\infty$, $\sigma \in \mathcal{P}(S)$, and $\chi_{[1\infty]} = (\chi_1, \chi_2, ...) \in (\mathcal{K}(S,X))^\infty$, the sequence $\{v^t(s, \xi_{[1t]}, \sigma, \chi_{[1t]}) \mid t = 0, 1, ...\}$ is thus Cauchy and has a limit point $v^\infty(s, \xi_{[1\infty]}, \sigma, \chi_{[1\infty]})$. The latter is the total discounted expected payoff a player can obtain in the game Γ^∞, when he starts at state s and adopts action plan $\xi_{[1\infty]}$, while all other players form initial pre-action environment σ and act according to $\chi_{[1\infty]}$.

A pre-action environment $\sigma \in \mathcal{P}(S)$ is said to be associated with $\chi \in \mathcal{K}(S,X)$ when

$$\sigma = T(\chi) \circ \sigma. \tag{35}$$

That is, we let environment σ be associated with action plan χ when the former is invariant under the one-period transition when all players adhere to the latter. For $\chi \in \mathcal{K}(S,X)$, we use χ^∞ to represent the stationary policy profile $(\chi, \chi, ...) \in (\mathcal{K}(S,X))^\infty$ that players are to adopt in all periods $t = 1, 2, ...$

We deem one-time action plan $\chi \in \mathcal{K}(S,X)$ a stationary Markov equilibrium for the nonatomic game Γ^∞, when there exists a $\sigma \in \mathcal{P}(S)$ that is associated with the given χ, so that for every one-time unilateral deviation $\xi \in \mathcal{K}(S,X)$,

$$v^\infty(s, \chi^\infty, \sigma, \chi^\infty) \geq v^\infty(s, (\xi, \chi^\infty), \sigma, \chi^\infty), \qquad \forall s \in S. \tag{36}$$

Therefore, a policy will be considered an equilibrium when it induces an invariant environment under whose sway the policy turns out to be a best response in the long run.

Now we move on to the n-player game Γ_n^∞ made out of the same S, X, $\bar{\alpha}$, $\tilde{f}$, and $\tilde{g}$. Similarly to the above, we let Γ_n^t be its n-player counterpart that terminates in period $t+1$. Now let $v_n^t(s_1, \xi_{[1t]}, \varepsilon_{s_{-1}}, \chi_{[1t]})$ be the total expected payoff player 1 can receive in game Γ_t^n, when he starts with state $s_1 \in S$ and adopts action plan $\xi_{[1t]} \in (\mathcal{K}(S,X))^t$ from period 1 to t,

while other players form initial empirical distribution $\varepsilon_{s_{-1}} = \varepsilon_{(s_2,...,s_n)} \in \mathcal{P}_{n-1}(S)$ and adopt policy $\chi_{[1t]} \in (\mathcal{K}(S,X))^t$ from 1 to t. As a terminal condition, we have $v_n^0(s_1, \varepsilon_{s_{-1}}) = 0$. For $t = 1, 2, ...$, it follows that

$$\begin{array}{c} v_n^t(s_1, \xi_{[1t]}, \varepsilon_{s_{-1}}, \chi_{[1t]}) = \int_X \xi_1(s_1|dx_1) \cdot \int_{X^{n-1}} \chi_1^{n-1}(s_{-1}|dx_{-1}) \cdot [\tilde{f}(s_1, x_1, \varepsilon_{s_{-1}x_{-1}}) \\ +\bar{\alpha} \cdot \int_{S^n} \tilde{g}^n(s, x|ds') \cdot v_n^{t-1}(s'_1, \xi_{[2t]}, \varepsilon_{s'_{-1}}, \chi_{[2t]})]. \end{array} \tag{37}$$

Using the terminal condition and (37), we can inductively show that

$$\mid v_n^{t+1}(s_1, \xi_{[1,t+1]}, \varepsilon_{s_{-1}}, \chi_{[1,t+1]}) - v_n^t(s_1, \xi_{[1t]}, \varepsilon_{s_{-1}}, \chi_{[1t]}) \mid \leq \bar{\alpha}^t \cdot \bar{f}. \tag{38}$$

Given $s_1 \in S$, $\xi_{[1\infty]} \in (\mathcal{K}(S,X))^\infty$, $\varepsilon_{s_{-1}} \in \mathcal{P}_{n-1}(S)$, and $\chi_{[1\infty]} \in (\mathcal{K}(S,X))^\infty$, the sequence $\{v_t^n(s_1, \xi_{[1t]}, \varepsilon_{s_{-1}}, \chi_{[1t]}) \mid t = 0, 1, ...\}$ is Cauchy and has a limit point $v_n^\infty(s_1, \xi_{[1\infty]}, \varepsilon_{s_{-1}}, \chi_{[1\infty]})$. The latter is the total discounted expected payoff a player can obtain in Γ_n^∞, when he starts at state s and adopts action plan $\xi_{[1\infty]}$, while all other players form the initial pre-action environment $\varepsilon_{s_{-1}}$ and act according to $\chi_{[1\infty]}$.

For the current setting, it should be noted that Assumptions 1 and 2 translate into the continuity in τ at an (s,x)-independent rate of, respectively, the transition kernel $\tilde{g}(s,x,\tau)$ and payoff function $\tilde{f}(s,x,\tau)$. We now present the main result for the stationary case.

Theorem 4 *Suppose $\chi \in \mathcal{K}(S,X)$ is a stationary Markov equilibrium for the stationary nonatomic game Γ^∞. Let $\hat{\pi}_{n-1} \in \mathcal{P}(S^{n-1})$ for each $n \in \mathbb{N} \setminus \{1\}$. Also suppose the sequence $\hat{\pi}_{n-1}$ asymptotically resembles the sequence σ^{n-1}, where σ is associated with χ in the equilibrium definitions (35) and (36). Then, χ^∞ would be asymptotically equilibrium for games Γ_n^∞ in an average sense. More specifically, for any $\epsilon > 0$ and large enough $n \in \mathbb{N}$,*

$$\int_{S^{n-1}} \hat{\pi}_{n-1}(ds_{-1}) \cdot v_n^\infty(s_1, \chi^\infty, \varepsilon_{s_{-1}}, \chi^\infty) \geq \int_{S^{n-1}} \hat{\pi}_{n-1}(ds_{-1}) \cdot v_n^\infty(s_1, \xi_{[1\infty]}, \varepsilon_{s_{-1}}, \chi^\infty) - \epsilon,$$

for any $s_1 \in S$ and $\xi_{[1\infty]} \in (\mathcal{K}(S,X))^\infty$.

Theorem 4 says that, players in a large finite stationary game will not regret much by adopting a stationary equilibrium for a correspondent stationary nonatomic game. The regret can be measured in an average sense, so long as the underlying other-player multi-state distribution $\hat{\pi}_{n-1}$ is close to an invariant σ associated with the NG equilibrium. Just as in Section 7, we can let $\hat{\pi}_{n-1} = \sigma^{n-1}$, indicating that players take a "lazy" approach in assessing other players' states. We leave discussion of other possibilities to Appendx E.

9 Implications of Main Results

9.1 Observation, Remembrance, and Coordination

Regarding Theorems 2 and 3, we note the following for $\bar{t}$-period games. A prominent feature of an NG equilibrium $\chi_{[1\bar{t}]} \in (\mathcal{K}(S,X))^{\bar{t}}$ is its insensitivity, at any period t, to a player's personal history $(s_{t'}, x_{t'}|t' = 1, 2, ..., t-1)$, historical data regarding other players, and the present information about other players' states. Independence of the first two factors has much to do with the Markovian setup of the game—neither $\tilde{f}_t$ nor $\tilde{g}_t$ depends on past history. But the more interesting independence of the latter two factors stems from players' common knowledge about the evolution of their environments. The $(\sigma_{t'} \otimes \chi_{t'}|t' = 1, 2, ..., t-1)$ portion of the history and the present information σ_t, both about other players, are determinable by (10) before the game is even played out.

For finite semi-anonymous games, however, information is gradually revealed and its perfection is not guaranteed. We can define space O_S and map $\tilde{o}_S : \mathcal{P}(S) \to O_S$ to represent a player's observatory power over his present pre-action environment immediately before actual play. Similarly, we can define space O_{SX} and map $\tilde{o}_{SX} : \mathcal{P}(S \times X) \to O_{SX}$ to represent his observatory power over the in-action environment just experienced. So that new information does not contradict old information and no information gets lost, we suppose function $\tilde{o}_S^{SX} : O_{SX} \to O_S$ exists, with $\tilde{o}_S^{SX}(\tilde{o}_{SX}(\tau)) = \tilde{o}_S(\tau|_S)$ for any $\tau \in \mathcal{P}(S \times X)$.

With these definitions, a player's decision in period t can be denoted by a map $\hat{\chi}_t : (S \times X \times O_{SX})^{t-1} \times O_S \times S \to \mathcal{P}(X)$. In the period, player 1's random decision rule can be written as $\hat{\chi}_t(\tilde{h}_t, \tilde{o}_S(\varepsilon_{s_{t,-1}}), s_{t1}|\cdot)$, where the history $\tilde{h}_t$ is expressible as

$$\tilde{h}_t = (s_{t'1}, x_{t'1}, \tilde{o}_{SX}(\varepsilon_{s_{t',-1}x_{t',-1}})|t' = 1, 2, ..., t-1), \tag{39}$$

$\tilde{o}_S(\varepsilon_{s_{t,-1}})$ is his observation of other players' status, and s_{t1} represents the player's own state. There is a whole spectrum in which O_S and $\tilde{o}_S$ can reside. When $O_S = \{0\}$ and $\tilde{o}_S(\cdot) = 0$, players are ignorant of others' states; when $O_S = \mathcal{P}(S)$ and $\tilde{o}_S$ is the identity map, every player is fully aware of his surrounding. Similarly, there are varieties of O_{SX}, $\tilde{o}_{SX}$, and $\tilde{o}_S^{SX}$.

Theorems 2 and 3, however, nullify the need to delve into the $(O_S, \tilde{o}_S, O_{SX}, \tilde{o}_{SX}, \tilde{o}_S^{SX})$-related details about finite games. They state that an equilibrium of the NG counterpart, which is necessarily both oblivious of the past history $\tilde{h}_t$ and blind to the present observation $\tilde{o}_S(\varepsilon_{s_{t,-1}})$, serves as a good approximate equilibrium for games with enough players. The absence of $\tilde{h}_t$ again has a Markovian explanation. On the other hand, the ability to

shake off $\tilde{o}_S(\varepsilon_{s_{t,-1}})$'s influence is very important, since this saves players the efforts to gather information about their surroundings.

Regarding Theorem 4, we note the following. Each of our finite stationary games is a discounted stochastic game. For an n-player version of the latter game in which players have full knowledge of others' states, equilibria are hard to compute and for their implementation, require high degrees of coordination among players; see Solan [30]. These equilibria come from the space $(2^{\mathbb{R}^n})^{S^n} \times ((\mathbb{R}^n)^{X^n \times S^n})^{S^n \times \mathbb{R}^n}$; whereas, our NG equilibria come from $\mathbb{R}^{S \times X}$. Meanwhile, the discounted stochastic game one faces in real life is often semi-anonymous; see, e.g., examples listed in Jovanoic and Rosenthal [16]. For such a game, Theorem 4 has shown that a much easier path can be taken in order to coordinate player behavior under an ϵ-sized compromise. If players all agree to exercise a corresponding NG equilibrium, the typical player 1 has only to respond to his own state s_{t1} without giving up too much.

9.2 Sources of NG Equilibria

To further buttress the claim that studying the idealistic NGs can help with the understanding and execution of messier finite games faced in real life, we demonstrate that NG equilibria, meeting criteria (23) and (24) for the transient case and (35) and (36) for the stationary case, can be obtained relatively easily.

First, we concentrate on the transient case studied in Sections 3 to 6. From (22),

$$v_t(s_t, (\xi_t, \chi_{[t+1,\bar{t}]}), \sigma_t, \chi_{[t\bar{t}]}) = \int_X \xi_t(dy) \cdot v_t(s_t, (\delta_y, \chi_{[t+1,\bar{t}]}), \sigma_t, \chi_{[t\bar{t}]}). \tag{40}$$

Hence,

$$\sup_{\xi_t \in \mathcal{K}(S,X)} v_t(s_t, (\xi_t, \chi_{[t+1,\bar{t}]}), \sigma_t, \chi_{[t\bar{t}]}) = \sup_{y \in X} v_t(s_t, (\delta_y, \chi_{[t+1,\bar{t}]}), \sigma_t, \chi_{[t\bar{t}]}). \tag{41}$$

So the equilibrium criterion (23) conveniently used by us for the $\bar{t}$-period case is equivalent to, for every $t = 1, 2, ..., \bar{t}$,

$$\chi_t(s_t | \tilde{X}_t(s_t, \sigma_t, \chi_{[t\bar{t}]})) = 1, \qquad \forall s_t \in S, \tag{42}$$

where

$$\tilde{X}_t(s_t, \sigma_t, \chi_{[t\bar{t}]}) = \{x \in X | v_t(s_t, (\delta_x, \chi_{[t+1,\bar{t}]}), \sigma_t, \chi_{[t\bar{t}]}) = \sup_{y \in X} v_t(s_t, (\delta_y, \chi_{[t+1,\bar{t}]}), \sigma_t, \chi_{[t\bar{t}]})\}, \tag{43}$$

and σ_t is defined through (24).

The form consisting of (42) and (43) is fairly close to the distributional-equilibrium concept used in NG literature, such as Mas-Colell [24] and Jovanovic and Rosenthal [16]. A distributional equilibrium is an in-action environment sequence $\tau_{[1\bar{t}]} = (\tau_t | t = 1, 2, ..., \bar{t}) \in (\mathcal{P}(S \times X))^{\bar{t}}$ which satisfies $\tau_t(\tilde{U}_t(\tau_{[t\bar{t}]})) = 1$ for each $t = 1, 2, ..., \bar{t}$. Here, $\tilde{U}_t(\tau_{[t\bar{t}]}) = \{(s, x) \in S \times X | v'_t(s, x, \tau_{[t\bar{t}]}) = \sup_{y \in X} v'_t(s, y, \tau_{[t\bar{t}]})\}$, and $v'_t(s, y, \tau_{[t\bar{t}]})$ is a player's payoff when he starts period t with state s and action y, but other players in all periods and he himself in later periods act according to $\tau_{[t\bar{t}]}$; corresponding to (24), the distributional equilibrium also satisfies $\tau_1|_S = \sigma_1$ and $\tau_t|_S = \tau_{t-1} \odot \tilde{g}_{t-1}(\cdot, \cdot, \tau_{t-1})$ for $t = 2, 3, ..., \bar{t}$. According to Jovanovic and Rosenthal [16] (Theorem 1), such an equilibrium $\tau_{[1\bar{t}]}$ would exist when S and X are compact, each payoff $\tilde{f}_t$ is bounded and continuous in all arguments, and each transition kernel $\tilde{g}_t$ is continuous in all arguments.

When an equilibrium $\chi_{[1\bar{t}]}$ in our conditional sense exists, we can construct a distributional equilibrium $\tau_{[1\bar{t}]}$ by resorting iteratively to $\tau_t = \sigma_t \otimes \chi_t$ and $\sigma_{t+1} = T_t(\chi_t) \circ \sigma_t$ for $t = 1, 2, ..., \bar{t}$. Conversely, when the latter distributional equilibrium $\tau_{[1\bar{t}]}$ is available, we can nearly get a conditional equilibrium $\chi_{[1\bar{t}]}$ back. For each $t = 1, 2, ..., \bar{t}$, according to Duffie, Geanakoplos, Mas-Colell, and McLennan [9] (p. 751), we can identify a $\chi_t \in \mathcal{K}(S, X)$, which also passes as a measurable map from S to $\mathcal{P}(X)$, that satisfies $\tau_t = \tau_t|_S \otimes \chi_t$. Thus, we will be able to construct $\chi_{[1\bar{t}]}$ consecutively from χ_1 up to $\chi_{\bar{t}}$. But even then, $\chi_{[t\bar{t}]}$ along with $\sigma_t = \tau_t|_S$ would satisfy (42) only for $\tau_t|_S$-almost every s_t, but not necessarily every $s_t \in S$. For instance, we can suppose $S = \{\bar{s}_1, \bar{s}_2, ...\}$. At each t, the constructed $\chi_{[1\bar{t}]}$ could guarantee (42) for those $\bar{s}_i$'s with $(\tau_t|_S)(\bar{s}_i) > 0$ but not those with $(\tau_t|_S)(\bar{s}_i) = 0$. On the other hand, a conditional equilibrium $\chi_{[1\bar{t}]}$ can be obtained directly; see Appendix F.1 for details.

When it comes to the stationary case examined in Section 8, we make parallel developments. Here the property corresponding to (36) is

$$\chi(s | \tilde{X}_\infty(s, \sigma, \chi)) = 1, \qquad \forall s \in S, \tag{44}$$

where

$$\tilde{X}_\infty(s, \sigma, \chi) = \{x \in X | v^\infty(s, (\delta_x, \chi^\infty), \sigma, \chi^\infty) = \sup_{y \in X} v^\infty(s, (\delta_y, \chi^\infty), \sigma, \chi^\infty)\}, \tag{45}$$

and σ satisfies (35). Again, the existence of a related distributional equilibrium $\tau \in \mathcal{P}(S \times X)$ is known under quite general conditions; see, e.g., Jovanovic and Rosenthal [16] (Theorem 2). However, an equilibrium τ does not exactly lead to a conditional equilibrium χ. So once more we focus on a direct approach for the stationary case; see Appendix F.2.

10 Concluding Remarks

Under a common action plan, we have shown that environments faced by players in multi-period large finite games would stay close to those of their NG counterparts. For transient and stationary settings, our results reveal that an NG equilibrium, necessarily both oblivious of past history and blind to present status of other players, could serve as a good approximate equilibrium in large finite games. We reckon that the discreteness requirement on both the state and action spaces can be frustrating in some circumstances. Besides the relaxation of the aforementioned restriction, future research can also look into the issue of converge rate.

Acknowledgments

This research was supported in part by National Science Foundation Grant CMMI-0854803, as well as National Natural Science Foundation of China Grants 11371273 and 71502015.

Appendices

A Concepts and Rudimentary Lemmas

Recall that ρ_A stands for the Prohorov metric for the space of distributions $\mathcal{P}(A)$.

Lemma 1 *Let A be a separable metric space. Then, for any $n \in \mathbb{N}$ and $a, a' \in A^n$,*

$$\rho_A(\varepsilon_a, \varepsilon_{a'}) \leq \max_{m=1}^{n} d_A(a_m, a'_m).$$

Proof: Let $\epsilon = \max_{m=1}^{n} d_A(a_m, a'_m)$. For any $A' \in \mathcal{B}(A)$, the key observation is that

$$\delta_{a'_m}((A')^\epsilon) \geq \delta_{a_m}(A'). \tag{A.1}$$

Then,

$$\varepsilon_{a'}((A')^\epsilon) = \frac{\sum_{m=1}^{n} \delta_{a'_m}((A')^\epsilon)}{n} \geq \frac{\sum_{m=1}^{n} \delta_{a_m}(A')}{n} = \varepsilon_a(A'). \tag{A.2}$$

Thus, $\rho_A(\varepsilon_a, \varepsilon_{a'}) \leq \epsilon$. ∎

According to Parthasarathy [26] (Theorem II.7.1), the strong law of large numbers applies to the empirical distribution under the weak topology, and hence under the Prohorov metric. In the following, we state its weak version.

Lemma 2 *Let separable metric space A and distribution $p \in \mathcal{P}(A)$ be given. Then, for any $\epsilon > 0$, as long as n is large enough,*

$$p^n(\{a \in A^n \mid \rho_A(\varepsilon_a, p) < \epsilon\}) > 1 - \epsilon.$$

Due to the inequality of Dvoretzky, Kiefer, and Wolfolwitz [10], the above convergence is uniform for certain A's. The inequality infers that, when A is $\mathbb{R}$ or countable,

$$p^n(\{a \in A^n \mid \rho_A(\varepsilon_a, p) \leq \epsilon\}) > 1 - 2e^{-2n\epsilon^2}, \qquad \forall \epsilon > 0.$$

When n is greater than $\ln(3/\epsilon)/(2\epsilon^2)$, a number independent of $p \in \mathcal{P}(A)$, the above would entail the inequality in Lemma 2. Thus, we have the following.

Lemma 3 *When A is the real line $\mathbb{R}$ or countable, the convergence expressed in Lemma 2 is uniform. Namely, a lower bound could be identified so that every n above it would realize the inequality in the lemma for every $p \in \mathcal{P}(A)$.*

For separable metric space A, point $a \in A$, and the $(n-1)$-point empirical distribution $p \in \mathcal{P}_{n-1}(A)$, we use $(a,p)_n$ to represent the member of $\mathcal{P}_n(A)$ that has an additional $1/n$ weight on the point a, but with probability masses in p being reduced to $(n-1)/n$ times of their original values. For $a \in A^n$ and $m = 1, ..., n$, we have $(a_m, \varepsilon_{a_{-m}})_n = \varepsilon_a$. Concerning the Prohorov metric, we have also a simple but useful observation.

Lemma 4 *Let A be a separable metric space. Then, for any $n \in \mathbb{N} \setminus \{1\}$, $a \in A$, and $p \in \mathcal{P}_{n-1}(A)$,*

$$\rho_A((a,p)_n, p) \leq \frac{1}{n}.$$

Proof: Let $A' \in \mathcal{B}(A)$ be chosen. Then $p(A') = (m-1)/(n-1)$ for some $m = 1, 2, ..., n$. If $a \notin A'$, then $(a,p)_n(A') = (m-1)/n$ and hence

$$(a,p)_n(A') \leq p(A') \leq (a,p)_n(A') + \frac{1}{n}. \tag{A.3}$$

If $a \in A'$, then $(a,p)_n(A') = m/n$ and hence

$$(a,p)_n(A') - \frac{1}{n} \leq p(A') \leq (a,p)_n(A'). \tag{A.4}$$

Therefore, it is always true that

$$\mid (a,p)_n(A') - p(A') \mid \leq \frac{1}{n}. \tag{A.5}$$

Due to the nature of the Prohorov metric, we have

$$\rho_A((a,p)_n, p) \leq \frac{1}{n}. \tag{A.6}$$

We have thus completed the proof. ∎

For the notion of asymptotic resemblance introduced in Definition 1, we have that it is preserved under certain projections and expansions.

Lemma 5 *Let A be a separable metric space. Also, $q_n \in \mathcal{P}(A^n)$ for every $n \in \mathbb{N}$ and $p \in \mathcal{P}(A)$. Suppose the sequence q_n asymptotically resembles the sequence p^n. Then, the sequence $q_n|_{A^{n-1}}$ will asymptotically resemble the sequence p^{n-1}.*

Proof: For any $\epsilon > 0$, due to the asymptotic resemblance of the sequence q_n to the sequence p^n, we have, for n large enough,

$$q_n(A'_n) > 1 - \epsilon, \tag{A.7}$$

where

$$A'_n = \{a \in A^n | \rho_A(\varepsilon_a, p) < \epsilon\}. \tag{A.8}$$

By Lemma 4, we have

$$\rho_A(\varepsilon_a, \varepsilon_{a_{-1}}) \leq \frac{1}{n}, \qquad \forall a \in A^n. \tag{A.9}$$

Hence, for large enough n,

$$A'_n \subseteq A \times A''_{n-1}, \tag{A.10}$$

where

$$A''_{n-1} = \{a_{-1} \in A^{n-1} | \rho_A(\varepsilon_{a_{-1}}, p) < 2\epsilon\}. \tag{A.11}$$

But by (A.7), this means that

$$(q_n|_{A^{n-1}})(A''_{n-1}) = q_n(A \times A''_{n-1}) \geq q_n(A'_n) > 1 - \epsilon. \tag{A.12}$$

That is, $q_n|_{A^{n-1}}$ asymptotically resembles p^{n-1}. ∎

Lemma 6 *Let A be a separable metric space. Also, $q_n \in \mathcal{P}(A^n)$ for every $n \in \mathbb{N}$ and $p, p' \in \mathcal{P}(A)$. Suppose the sequence q_n asymptotically resembles the sequence p^n. Then, the sequence $p' \times q_{n-1}$ will asymptotically resemble the sequence p^n as well.*

Proof: For any $\epsilon > 0$, due to the asymptotic resemblance of the sequence q_n to the sequence p^n, we have, for n large enough,

$$q_{n-1}(A'_{n-1}) > 1 - \epsilon, \tag{A.13}$$

where

$$A'_{n-1} = \{a \in A^{n-1} | \rho_A(\varepsilon_a, p) < \epsilon\}. \tag{A.14}$$

By Lemma 4, we have

$$\rho_A(\varepsilon_{(a_1,a)}, \varepsilon_a) \leq \frac{1}{n}, \qquad \forall a_1 \in A,\ a \in A^{n-1}. \tag{A.15}$$

Hence, for large enough n,

$$A \times A'_{n-1} \subseteq A''_n, \tag{A.16}$$

where

$$A''_n = \{a \in A^n | \rho_A(\varepsilon_a, p) < 2\epsilon\}. \tag{A.17}$$

But by (A.13), this means that

$$(p' \times q_{n-1})(A''_n) \geq (p' \times q_{n-1})(A \times A'_{n-1}) = p'(A) \times q_{n-1}(A'_{n-1}) > 1 - \epsilon. \tag{A.18}$$

That is, $p' \times q_{n-1}$ asymptotically resembles p^n. ∎

Lemma 7 *Let A and B be separable metric spaces. Also, $q_n \in \mathcal{P}(A^n \times B^n)$ for every $n \in \mathbb{N}$ and $p \in \mathcal{P}(A \times B)$. Suppose the sequence q_n asymptotically resembles the sequence p^n. Then, the sequence $q_n|_{A^n}$ will asymptotically resemble the sequence $(p|_A)^n$.*

Proof: For any $\epsilon > 0$, due to the asymptotic resemblance of the sequence q_n to the sequence p^n, we have, for n large enough,

$$q_n(C'_n) > 1 - \epsilon, \tag{A.19}$$

where

$$C'_n = \{c = (a, b) \in A^n \times B^n | \rho_{A\times B}(\varepsilon_c, p) < \epsilon\}. \tag{A.20}$$

But by (87) of Yang [33],

$$\rho_A(\varepsilon_a, p|_A) = \rho_A(\varepsilon_c|_A, p|_A) \leq \rho_{A\times B}(\varepsilon_c, p), \qquad \forall c = (a,b) \in C'_n. \tag{A.21}$$

Hence,

$$C'_n \subseteq A'_n \times B^n, \tag{A.22}$$

where

$$A'_n = \{a \in A^n | \rho_A(\varepsilon_a, p|_A) < \epsilon\}. \tag{A.23}$$

Combining (A.19) and (A.22), we can obtain

$$(q_n|_{A^n})(A'_n) = q_n(A'_n \times B^n) \geq q_n(C'_n) > 1 - \epsilon. \tag{A.24}$$

This indicates that $q_n|_{A^n}$ asymptotically resembles $(p|_A)^n$. ∎

B Proofs of Section 5

Proof of Proposition 1: We first prove (i). Fix some $\epsilon \in (0,1)$. Due to the countability of S , we can identify some I of its points $\bar{s}_1, \bar{s}_2, ..., \bar{s}_I$, so that each $\sigma(\{\bar{s}_i\}) > 0$ and

$$\sum_{i=1}^{I} \sigma(\{\bar{s}_i\}) > 1 - \epsilon. \tag{B.1}$$

For convenience, let $\bar{S}' = \{\bar{s}_1, \bar{s}_2, ..., \bar{s}_I\}$ and $\bar{S}'' = S \setminus \bar{S}'$.

Since S is discrete, the distance $d_S(\bar{S}', \bar{S}'') = \inf_{s'\in\bar{S}', s''\in\bar{S}''} d_S(s', s'') > 0$. For $i, j = 1, 2, ..., I$, let us use d_{ij} for $d_S(\bar{s}_i, \bar{s}_j)$ and σ_i for $\sigma(\{\bar{s}_i\})$. Now define

$$\delta = \frac{\epsilon}{I} \wedge d_S(\bar{S}', \bar{S}'') \wedge (\min_{i\neq j} d_{ij}) \wedge (\min_i \frac{\sigma_i}{2}), \tag{B.2}$$

which is still strictly positive. In this paper, we use $a \wedge b$ to stand for $\min\{a, b\}$ and $a \vee b$ to stand for $\max\{a, b\}$.

For any $n \in \mathbb{N}$, define $S'_n \in \mathcal{B}(S^n)$ so that

$$S'_n = \{s \in S^n | \rho_S(\varepsilon_s, \sigma) < \delta\}. \tag{B.3}$$

By the hypothesis that π_n asymptotically resembles σ^n, we can ensure

$$\pi_n(S'_n) > 1 - \frac{\epsilon}{2}, \tag{B.4}$$

by making n large enough.

Consider any such n, as well as any $s = (s_1, s_2, ..., s_n) \in S'_n$ and $i = 1, 2, ..., I$. It follows from $\delta \leq d_S(\bar{S}', \bar{S}'') \wedge (\min_{i \neq j} d_{ij})$ that $(\{\bar{s}_i\})^\delta$, whose meaning comes from (**??**), is still $\{\bar{s}_i\}$ itself. Now by (B.3),

$$\varepsilon_s(\{\bar{s}_i\}) < \sigma((\{\bar{s}_i\})^\delta) + \delta = \sigma_i + \delta, \tag{B.5}$$

and

$$\begin{aligned} \varepsilon_s(\{\bar{s}_i\}) &= 1 - \varepsilon_s(\{\bar{s}_j | j \neq i\} \cup \bar{S}'') > 1 - \sigma((\{\bar{s}_j | j \neq i\} \cup \bar{S}'')^\delta) - \delta \\ &= 1 - \sigma(\{\bar{s}_j | j \neq i\} \cup \bar{S}'') - \delta = \sigma_i - \delta, \end{aligned} \tag{B.6}$$

which is still above $\delta > 0$ by the fact that $\delta \leq \min_i \sigma_i/2$. For convenience, let $n_i(s) = n \cdot \varepsilon_s(\{\bar{s}_i\})$, the number of components s_m of s that happen to be $\bar{s}_i$. Now we know that $n_i(s)$ is above $n\delta$ for every $s \in S'_n$ and $i = 1, 2, ..., I$.

On the other hand, by Lemma 2, there exists some $\underline{n}_i$ for each $i = 1, 2, ..., I$, so that when $n_i > \underline{n}_i$,

$$(\chi(\bar{s}_i))^{n_i}(X'_{in_i}) > 1 - \frac{\epsilon}{2I}, \tag{B.7}$$

where

$$X'_{in_i} = \{x \in X^{n_i} | \rho_X(\varepsilon_x, \chi(\bar{s}_i)) < \delta\}. \tag{B.8}$$

Since $\delta > 0$, we can ensure that $n\delta$ and hence $n_i(s)$ is above $\underline{n}_i$ for every $i = 1, 2, ..., I$ by letting n be large enough.

Fix a big n that facilitates both (B.4) and (B.7). For any $(s, x) \in S^n \times X^n$, let $\tilde{x}_i(s, x)$ be the $n_i(s)$-long vector of x_m's whose corresponding s_m's happen to be $\bar{s}_i$:

$$\tilde{x}_i(s, x) = (x_m | m = 1, 2, ..., n \text{ but with } s_m = \bar{s}_i) \in X^{n_i(s)}. \tag{B.9}$$

Define $U'_n \in \mathcal{B}(S^n \times X^n)$, so that

$$U'_n = \{(s, x) \in S^n \times X^n | s \in S'_n \text{ and } \tilde{x}_i(s, x) \in X'_{in_i(s)} \text{ for each } i = 1, 2, ..., I\}. \tag{B.10}$$

By (16), (B.4), and (B.7), we have

$$(\pi_n \otimes \chi^n)(U'_n) = \int_{S'_n} \pi_n(ds) \cdot \prod_{i=1}^{I} (\chi(\bar{s}_i))^{n_i(s)}(X'_{in_i(s)}) > (1 - \frac{\epsilon}{2}) \cdot (1 - \frac{\epsilon}{2I})^I > 1 - \epsilon. \tag{B.11}$$

For any (s, x) in U'_n, let us examine how close $\varepsilon_{sx} = \varepsilon_{((s_1,x_1),...,(s_n,x_n))}$ is to $\sigma \otimes \chi$. Recall that $S = \{\bar{s}_1, \bar{s}_2, ..., \bar{s}_I\} \cup \bar{S}''$. So for any $U' \in \mathcal{B}(S \times X)$,

$$U' = (\bigcup_{i=1}^{I} \{\bar{s}_i\} \times X'_i) \bigcup U'', \tag{B.12}$$

where $X_i' \in \mathcal{B}(X)$ for $i = 1, 2, ..., I$, while U'' is such that $s'' \in \bar{S}''$ for any $(s'', x'') \in U''$. Note again that $\delta \leq d_S(\bar{S}', \bar{S}'') \wedge \min_{i \neq j} d_{ij}$. When we take $d_{S \times X}$ to mean $d_{S \times X}((s', x'), (s'', x'')) = d_S(s', s'') \vee d_X(x', x'')$, (B.12) would lead to

$$\bigcup_{i=1}^{I} \{\bar{s}_i\} \times (X_i')^\delta \subseteq (U')^\delta. \tag{B.13}$$

Now from (B.5) and (B.8),

$$\begin{aligned} \varepsilon_{sx}(\{\bar{s}_i\} \times X_i') &= \varepsilon_s(\{\bar{s}_i\}) \cdot \varepsilon_{\tilde{x}_i(s,x)}(X_i') < (\sigma_i + \delta) \cdot [\chi(\bar{s}_i|(X_i')^\delta) + \delta] \\ &\leq (\sigma \otimes \chi)(\{\bar{s}_i\} \times (X_i')^\delta) + 2\delta + \delta^2 < (\sigma \otimes \chi)(\{\bar{s}_i\} \times (X_i')^\delta) + 3\delta, \end{aligned} \tag{B.14}$$

where the last inequality is due to our choice that $\delta \leq \epsilon / I < 1$. Meanwhile,

$$\varepsilon_{sx}(U'') \leq \varepsilon_{sx}(\bar{S}'' \times X) = \varepsilon_s(\bar{S}'') = 1 - \sum_{i=1}^{I} \varepsilon_s(\{\bar{s}_i\}) < 1 - \sum_{i=1}^{I} \sigma_i + I\delta < \epsilon + I\delta, \tag{B.15}$$

where the second-to-last inequality is due to (B.6) and the last one is due to (B.1). Combine (B.12) to (B.15), and we can obtain

$$\varepsilon_{sx}(U') < (\sigma \otimes \chi)((U')^\delta) + \epsilon + 4I\delta. \tag{B.16}$$

Thus,

$$\rho_{S \times X}(\varepsilon_{sx}, \sigma \otimes \chi) < \epsilon + 4I\delta \leq 5\epsilon, \tag{B.17}$$

where the last inequality comes from our choice that $\delta \leq \epsilon / I$. Since (B.11) and (B.17) are to occur at any n that is large enough, we see that (i) is true.

We then prove (ii). For convenience, we denote $S \times X$ by U, $\sigma \otimes \chi$ by τ, and for each $n \in \mathbb{N}$, $\pi_n \otimes \chi^n$ by ν_n. From (i), we have the sequence ν_n asymptotically resembling the sequence τ^n.

Fix some $\epsilon \in (0, 1)$. Due to the countability of S and X , and hence that of U, we can identify some J points $\bar{u}_1, \bar{u}_2, ..., \bar{u}_J$, so that each $\tau(\{\bar{u}_j\}) > 0$ and

$$\sum_{j=1}^{J} \tau(\{\bar{u}_j\}) > 1 - \epsilon. \tag{B.18}$$

For convenience, let $\bar{U}' = \{\bar{u}_1, \bar{u}_2, ..., \bar{u}_J\}$ and $\bar{U}'' = U \setminus \bar{U}'$.

As S and X are both discrete, so U is discrete as well. Thence, the distance $d_U(\bar{U}', \bar{U}'') = \inf_{u' \in \bar{U}', u'' \in \bar{U}''} d_U(u', u'') > 0$. For $j, k = 1, 2, ..., J$, let us use d'_{jk} for $d_U(\bar{u}_j, \bar{u}_k)$ and τ_j for $\tau(\{\bar{u}_j\})$. Now define

$$\delta = \frac{\epsilon}{J} \wedge d_U(\bar{U}', \bar{U}'') \wedge (\min_{j \neq k} d'_{jk}) \wedge (\min_j \frac{\tau_j}{2}), \tag{B.19}$$

which is still strictly positive.

For any $n \in \mathbb{N}$, define $U'_n \in \mathcal{B}(U^n)$ so that

$$U'_n = \{u \in U^n | \rho_U(\varepsilon_u, \tau) \bigvee [2 \cdot \sup_{u' \in U} \max_{m=1}^{n} \rho_S(g(u', \varepsilon_{u_{-m}}), g(u', \tau))] < \delta\}. \tag{B.20}$$

By (i) that ν_n asymptotically resembles τ^n, the hypothesis that $g(u, \cdot)$ is continuous at a u-independent rate, and Lemma 4, we can ensure

$$\nu_n(U'_n) > 1 - \frac{\epsilon}{2}, \tag{B.21}$$

by making n large enough,

Consider any such n, as well as any $u = (u_1, u_2, ..., u_n) \in U'_n$ and $j = 1, 2, ..., J$. It follows from $\delta \leq d_U(\bar{U}', \bar{U}'') \wedge (\min_{j \neq k} d'_{jk})$ that $(\{\bar{u}_j\})^\delta$ is still $\{\bar{u}_j\}$ itself. Now by (B.20),

$$\varepsilon_u(\{\bar{u}_j\}) < \tau((\{\bar{u}_j\})^\delta) + \delta = \tau_j + \delta, \tag{B.22}$$

and

$$\varepsilon_u(\{\bar{u}_j\}) > 1 - \tau((\{\bar{u}_k | k \neq j\} \cup \bar{U}'')^\delta) - \delta = \tau_j - \delta, \tag{B.23}$$

which is still above $\delta > 0$ by the fact that $\delta \leq \min_j \tau_j / 2$. For convenience, let $n'_j(u) = n \cdot \varepsilon_u(\{\bar{u}_j\})$. Now we know that $n'_j(u)$ is above $\lfloor n\delta \rfloor$ for every $j = 1, 2, ..., J$.

Due to the countability of U and Lemma 3 on the uniform Glivenko-Cantelli property, there exists some $\underline{n}'$, independent of both j and u, such that when every $n'_j(u) > \underline{n}'$,

$$(g(\bar{u}_j, \varepsilon_{u \setminus \bar{u}_j}))^{n'_j(u)}(S''_{jn'_j(u)}(u)) > 1 - \frac{\epsilon}{2J}, \qquad \forall j = 1, 2, ..., J, \tag{B.24}$$

where every $u \setminus \bar{u}_j$ is the $(n-1)$-long vector that is almost identical to u but with only $n'_j(u) - 1$ components equal to $\bar{u}_j$, and

$$S''_{jn'}(u') = \{s \in S^{n'} | \rho_S(\varepsilon_s, g(\bar{u}_j, \varepsilon_{u' \setminus \bar{u}_j})) < \frac{\delta}{2}\}. \tag{B.25}$$

But in light of (B.20), we can really guarantee that

$$(g(\bar{u}_j, \varepsilon_{u \setminus \bar{u}_j}))^{n'_j(u)}(S'_{jn'_j(u)}) > 1 - \frac{\epsilon}{2J}, \qquad \forall j = 1, 2, ..., J, \tag{B.26}$$

where

$$S'_{jn'} = \{s \in S^{n'} | \rho_S(\varepsilon_s, g(\bar{u}_j, \tau)) < \delta\}. \tag{B.27}$$

Since $\delta > 0$, we can ensure that $\lfloor n\delta \rfloor$ and hence $n'_j(u)$ is above $\underline{n}'$ for every $j = 1, 2, ..., J$ by letting n be large enough.

Fix a big n that facilitates both (B.21) and (B.26). For any $(u,s) \in U^n \times S^n$, let $\tilde{s}_j(u,s)$ be the $n'_j(u)$-long vector of s_m's whose corresponding u_m's happen to be $\bar{u}_j$:

$$\tilde{s}_j(u,s) = (s_m | m = 1,2,...,n \text{ but with } u_m = \bar{u}_j) \in S^{n'_j(u)}. \tag{B.28}$$

Define $V'_n \in \mathcal{B}(U^n \times S^n)$, so that

$$V'_n = \{(u,s) \in U^n \times S^n | u \in U'_n \text{ and } \tilde{s}_j(u,s) \in S'_{jn'_j(u)} \text{ for each } j = 1,2,...,J\}. \tag{B.29}$$

Let us follow the same logic as used from (B.11) to (B.17) in the proof of (i), with appropriate substitutions, such as J for I, U for S, S for X, ν_n for π_n, τ for σ, $g(\cdot,\cdot,\tau)$ for χ, g^n for χ^n, V'_n for U'_n, (B.21) for (B.4), and (B.26) for (B.7). We can then derive that

$$(\nu_n \otimes g^n)(V'_n) > 1 - \epsilon, \tag{B.30}$$

whereas, for any (u,s) in V'_n,

$$\rho_{U \times S}(\varepsilon_{us}, \tau \otimes g(\cdot,\cdot,\tau)) < 5\epsilon. \tag{B.31}$$

Since (B.30) and (B.31) are to occur at any n that is large enough, we see that $\nu_n \otimes g^n$ would asymptotically resemble $(\tau \otimes g(\cdot,\cdot,\tau))^n$. Lemma 7 will then lead to the asymptotic resemblance of the sequence $\nu_n \odot g^n = (\nu_n \otimes g^n)|_{S^n}$ to the sequence $(\tau \odot g(\cdot,\cdot,\tau))^n = ((\tau \otimes g(\cdot,\cdot,\tau))|_S)^n$. Thus (ii) is true.

For (iii), denote the given (s,x) by u_1. By Lemma 4, we can make $\varepsilon_u = (u_1, \varepsilon_{u_{-1}})_n$ arbitrarily close to $\varepsilon_{u_{-1}}$ for any $u_{-1} = (u_2, u_3, ..., u_n) \in U^{n-1}$ by letting n be large enough. Hence, we can follow the proof of (ii) almost verbatim, with its (B.20) replaced by

$$U'_{n-1} = \{u_{-1} \in U^{n-1} | \rho_U(\varepsilon_u, \tau) \bigvee [2 \cdot \sup_{u' \in U} \max_{m=1}^{n} \rho_S(g(u', \varepsilon_{u_{-m}}), g(u',\tau))] < \delta\}, \tag{B.32}$$

its (B.21) replaced by

$$(\delta_{u_1} \times \nu_{n-1})(\{u_1\} \times U'_{n-1}) = \nu_{n-1}(U'_{n-1}) > 1 - \frac{\epsilon}{2}, \tag{B.33}$$

any choice of $u \in U^n$ replaced by $u_{-1} \in U^{n-1}$, and any choice of $u \in U'_n$ replaced by $u_{-1} \in U'_{n-1}$. ∎

Proof of Theorem 1: We prove by induction on t'.

First, note that $T_{[t,t-1]} \circ \sigma_t$ is merely σ_t itself. Hence, the claim is true for $t' = t$ because by the hypothesis, we do have π_{nt} asymptotically resembling $(T_{[t,t-1]} \circ \sigma_t)^n = {\sigma_t}^n$. Then,

for some $t' = t, t+1, ..., \bar{t}$, suppose the claim is true, that $\pi_{nt'} = \pi_{nt} \odot \Pi_{t''=t}^{t'-1}(\chi_{t''}^n \odot \tilde{g}_{t''}^n)$ asymptotically resembles $\sigma_{t'}^n = (T_{[t,t'-1]}(\chi_{[t,t'-1]}) \circ \sigma_t)^n$.

Assumption 1 on $\tilde{g}_{t'}(s, x, \tau)$'s equi-continuity in τ allows us to use part (ii) of Proposition 1. By it, we would have $\pi_{nt'} \odot \chi_{t'}^n \odot \tilde{g}_{t'}^n$ asymptotically resembling $(\sigma_{t'} \odot \chi_{t'} \odot \tilde{g}_t(\cdot, \cdot, \sigma_{t'} \otimes \chi_{t'}))^n$. Since the former is merely $\pi_{n,t'+1} = \pi_{nt} \odot \Pi_{t''=t}^{t'}(\chi_{t''}^n \odot \tilde{g}_{t''}^n)$ and the latter is ${\sigma_{t'}}_{+1}^n = (T_{[tt']}(\chi_{[tt']}) \circ \sigma_t)^n$, we have thus proved the claim for $t'+1$.

The induction process is now complete. ∎

C Proofs of Section 6

Proof of Proposition 2: Let us prove by induction on t. By (21) and (25), the desired result is true for $t = \bar{t} + 1$.

At some $t = \bar{t}, \bar{t}-1, ..., 1$, suppose for any σ_{t+1} and any sequence $\hat{\pi}_{n-1,t+1}$ that asymptotically resembles σ_{t+1}^{n-1}, the sequence $\int_{S^{n-1}} \hat{\pi}_{n-1,t+1}(ds_{t+1,-1}) \cdot v_{n,t+1}(s_{t+1,1}, \xi_{[t+1,\bar{t}]}, , \varepsilon_{s_{t+1,-1}}, \chi_{[t+1,\bar{t}]})$ converges to $v_{t+1}(s_{t+1,1}, \xi_{[t+1,\bar{t}]}, \sigma_{t+1}, \chi_{[t+1,\bar{t}]})$ at a rate independent of both $s_{t+1,1}$ and $\xi_{[t+1,\bar{t}]}$.

Now, given the sequence $\hat{\pi}_{n-1,t}$ that is known to asymptotically resemble ${\sigma_t}^{n-1}$, we are to show that $\int_{S^{n-1}} \hat{\pi}_{n-1,t}(ds_{t,-1}) \cdot v_{nt}(s_{t1}, \xi_{[t\bar{t}]}, \varepsilon_{s_{t,-1}}, \chi_{[t\bar{t}]})$ will converge to $v_t(s_{t1}, \xi_{[t\bar{t}]}, \sigma_t, \chi_{[t\bar{t}]})$ at a rate independent of both s_{t1} and $\xi_{[t\bar{t}]}$. For convenience, let $\sigma_{t+1} = T_t(\chi_t) \circ \sigma_t$.

Note that, by (22) and (26),

$$\begin{aligned} &\sup_{s_{t1} \in S, \xi_{[t\bar{t}]} \in (\mathcal{K}(S,X))^{\bar{t}-t+1}} \mid v_t(s_{t1}, \xi_{[t\bar{t}]}, \sigma_t, \chi_{[t\bar{t}]}) \\ &\qquad - \int_{S^{n-1}} \hat{\pi}_{n-1,t}(ds_{t,-1}) \cdot v_{nt}(s_{t1}, \xi_{[t\bar{t}]}, \varepsilon_{s_{t,-1}}, \chi_{[t\bar{t}]}) \mid \leq M_{n1} + M_{n2} + M_{n3}, \end{aligned} \tag{C.1}$$

where

$$\begin{aligned} M_{n1} = \sup_{(s_{t1}, x_{t1}) \in S \times X} &\int_{S^{n-1} \times X^{n-1}} (\hat{\pi}_{n-1,t} \otimes {\chi_t}^{n-1})(ds_{t,-1} \times dx_{t,-1}) \times \\ &\times \mid \tilde{f}_t(s_{t1}, x_{t1}, \sigma_t \otimes \chi_t) - \tilde{f}_t(s_{t1}, x_{t1}, \varepsilon_{s_{t,-1} x_{t,-1}}) \mid, \end{aligned} \tag{C.2}$$

$$\begin{aligned} M_{n2} = &\sup_{(s_{t1}, x_{t1}) \in S \times X, \xi_{[t+1,\bar{t}]} \in (\mathcal{K}(S,X))^{\bar{t}-t}} \int_S \tilde{g}_t(s_{t1}, x_{t1}, \sigma_t \otimes \chi_t | ds_{t+1,1}) \times \\ &\times \mid v_{t+1}(s_{t+1,1}, \xi_{[t+1,\bar{t}]}, \sigma_{t+1}, \chi_{[t+1,\bar{t}]}) - \int_{S^{n-1} \times X^{n-1}} (\hat{\pi}_{n-1,t} \otimes \chi_t^{n-1})(ds_{t,-1} \times dx_{t,-1}) \times \\ &\times \Pi_{m=2}^n \int_S \tilde{g}_t(s_{tm}, x_{tm}, \varepsilon_{s_{t,-m} x_{t,-m}} | ds_{t+1,m}) \cdot v_{n,t+1}(s_{t+1,1}, \xi_{[t+1,\bar{t}]}, \varepsilon_{s_{t+1,-1}}, \chi_{[t+1,\bar{t}]}) \mid, \end{aligned} \tag{C.3}$$

and

$$\begin{aligned} M_{n3} = &\sup_{(s_{t1}, x_{t1}) \in S \times X, \xi_{[t+1,\bar{t}]} \in (\mathcal{K}(S,X))^{\bar{t}-t}} \int_{S^{n-1} \times X^{n-1}} (\hat{\pi}_{n-1,t} \otimes {\chi_t}^{n-1})(ds_{t,-1} \times dx_{t,-1}) \times \\ &\times \mid [\int_S \tilde{g}_t(s_{t1}, x_{t1}, \sigma_t \otimes \chi_t | ds_{t+1,1}) - \int_S \tilde{g}_t(s_{t1}, x_{t1}, \varepsilon_{s_{t,-1} x_{t,-1}} | ds_{t+1,1})] \times \\ &\times \Pi_{m=2}^n \int_S \tilde{g}_t(s_{tm}, x_{tm}, \varepsilon_{s_{t,-m} x_{t,-m}} | ds_{t+1,m}) \cdot v_{n,t+1}(s_{t+1,1}, \xi_{[t+1,\bar{t}]}, \varepsilon_{s_{t+1,-1}}, \chi_{[t+1,\bar{t}]}) \mid . \end{aligned} \tag{C.4}$$

We now show that each of the above three terms can be made arbitrarily small by letting n be large enough.

For M_{n1}, define $\tilde{U}_{n-1}(\delta) \in \mathcal{B}(S^{n-1} \times X^{n-1})$ for every $\delta > 0$, so that

$$\tilde{U}_{n-1}(\delta) = \{(s_{t,-1}, x_{t,-1}) \in S^{n-1} \times X^{n-1} | \rho_{S\times X}(\varepsilon_{s_{t,-1}x_{t,-1}}, \sigma_t \otimes \chi_t) < \delta\}. \tag{C.5}$$

From (C.2), we know $M_{n1} \leq M_{n11}(\delta) + M_{n12}(\delta)$ for any $\delta > 0$, where

$$M_{n11}(\delta) = \sup_{(s_{t1},x_{t1})\in S\times X, (s_{t,-1},x_{t,-1})\in \tilde{U}_{n-1}(\delta)} | \tilde{f}_t(s_{t1}, x_{t1}, \sigma_t \otimes \chi_t) - \tilde{f}_t(s_{t1}, x_{t1}, \varepsilon_{s_{t,-1}x_{t,-1}}) |, \tag{C.6}$$

and

$$\begin{aligned} M_{n12}(\delta) = \sup_{(s_{t1},x_{t1})\in S\times X} \int_{(S^{n-1}\times X^{n-1})\setminus \tilde{U}_{n-1}(\delta)} (\hat{\pi}_{n-1,t} \otimes {\chi_t}^{n-1})(ds_{t,-1} \times dx_{t,-1}) \times \\ \times [| \tilde{f}_t(s_{t1}, x_{t1}, \sigma_t \otimes \chi_t) | + | \tilde{f}_t(s_{t1}, x_{t1}, \varepsilon_{s_{t,-1}x_{t,-1}}) |]. \end{aligned} \tag{C.7}$$

Because Assumption 2 says that $\tilde{f}_t(s, x, \tau)$ is continuous in τ at an (s, x)-independent rate, we can make $M_{n11}(\delta)$ arbitrarily small by letting δ be small enough. Meanwhile, by the asymptotic resemblance of the sequence $\hat{\pi}_{n-1,t}$ to the sequence ${\sigma_t}^{n-1}$ and part (i) of Proposition 1, we know that the sequence $\hat{\pi}_{n-1,t} \otimes {\chi_t}^{n-1}$ asymptotically resembles the sequence $(\sigma_t \otimes \chi_t)^{n-1}$. So the measure $(\hat{\pi}_{n-1,t} \otimes {\chi_t}^{n-1})((S^{n-1} \times X^{n-1}) \setminus \tilde{U}_{n-1}(\delta))$ can be made arbitrarily small at any δ by letting n be large enough. Since $\tilde{f}_t$ is bounded, this means that $M_{n12}(\delta)$ can be made arbitrarily small as well.

For M_{n2}, note the second integral in (C.3) can be understood as $\hat{\pi}_{n-1,t+1}(s_{t1}, x_{t1}|ds_{t+1,-1}) = \int_{S^{n-1}} \{[(\delta_{s_{t1}x_{t1}} \times ({\sigma_t}^{n-1} \otimes {\chi_t}^{n-1})) \odot \tilde{g}_t^n]|_{S^{n-1}}\}(ds_{t+1,-1})$. So we have

$$\begin{aligned} M_{n2} \leq \sup_{(s_{t1},x_{t1})\in S\times X, s_{t+1,1}\in S, \xi_{[t+1,\bar{t}]}\in (\mathcal{K}(S,X))^{\bar{t}-t}} | v_{t+1}(s_{t+1,1}, \xi_{[t+1,\bar{t}]}, \sigma_{t+1}, \chi_{[t+1,\bar{t}]}) \\ - \int_{S^{n-1}} \hat{\pi}_{n-1,t+1}(s_{t1}, x_{t1}|ds_{t+1,-1}) \cdot v_{n,t+1}(s_{t+1,1}, \xi_{[t+1,\bar{t}]}, \varepsilon_{s_{t+1,-1}}, \chi_{[t+1,\bar{t}]}) | . \end{aligned} \tag{C.8}$$

Meanwhile, Assumption 1 allows us to use part (iii) of Proposition 1. By the asymptotic resemblance of the sequence $\hat{\pi}_{n-1,t}$ to the sequence ${\sigma_t}^{n-1}$, part (iii) of Proposition 1, and Lemma 5, we know that the sequence $\hat{\pi}_{n-1,t+1}(s_{t1}, x_{t1})$ asymptotically resembles the sequence σ_{t+1}^{n-1} at an (s_{t1}, x_{t1})-independent rate. Then by the induction hypothesis where the convergence rate is also $(s_{t+1,1}, \xi_{[t+1,\bar{t}]})$-independent, we can conclude that M_{n2} can be made arbitrarily small by letting n be large enough.

For M_{n3}, define $V_n(s_{t1}, x_{t1}, \varepsilon_{s_{t,-1}x_{t,-1}}, \xi_{[t+1,\bar{t}]})$ so that

$$\begin{aligned} V_n(s_{t1}, x_{t1}, \varepsilon_{s_{t,-1}x_{t,-1}}, \xi_{[t+1,\bar{t}]}) =| [\int_S \tilde{g}_t(s_{t1}, x_{t1}, \sigma_t \otimes \chi_t|ds_{t+1,1}) \\ - \int_S \tilde{g}_t(s_{t1}, x_{t1}, \varepsilon_{s_{t,-1}x_{t,-1}}|ds_{t+1,1})] \cdot \Pi_{m=2}^n \tilde{g}_t(s_{tm}, x_{tm}, \varepsilon_{s_{t,-m}x_{t,-m}}|ds_{t+1,m}) \times \\ \times v_{n,t+1}(s_{t+1,1}, \xi_{[t+1,\bar{t}]}, \varepsilon_{s_{t+1,-1}}, \chi_{[t+1,\bar{t}]}) | . \end{aligned} \tag{C.9}$$

Then, (C.4) can be written as

$$M_{n3} = \sup_{(s_{t1},x_{t1})\in S\times X,\xi_{[t+1,\bar{t}]}\in(\mathcal{K}(S,X))^{\bar{t}-t}} \int_{S^{n-1}\times X^{n-1}} \times \\ \times(\hat{\pi}_{n-1,t}\otimes {\chi_t}^{n-1})(ds_{t,-1}\times dx_{t,-1})\cdot V_n(s_{t1},x_{t1},\varepsilon_{s_{t,-1}x_{t,-1}},\xi_{[t+1,\bar{t}]}). \tag{C.10}$$

Noting the definition of $\tilde{U}_{n-1}(\delta)$ in (C.5) for any $\delta>0$, we see that $M_{n3}\leq M_{n31}(\delta)+M_{n32}(\delta)$, where

$$M_{n31}(\delta) = \sup_{(s_{t1},x_{t1})\in S\times X,(s_{t,-1},x_{t,-1})\in\tilde{U}_{n-1}(\delta),\xi_{[t+1,\bar{t}]}\in(\mathcal{K}(S,X))^{\bar{t}-t}} \\ V_n(s_{t1},x_{t1},\varepsilon_{s_{t,-1}x_{t,-1}},\xi_{[t+1,\bar{t}]}), \tag{C.11}$$

and

$$M_{n32}(\delta) = \sup_{(s_{t1},x_{t1})\in S\times X,\xi_{[t+1,\bar{t}]}\in(\mathcal{K}(S,X))^{\bar{t}-t}} \int_{(s_{t,-1},x_{t,-1})\in(S^{n-1}\times X^{n-1})\setminus\tilde{U}_{n-1}(\delta)} \\ (\hat{\pi}_{n-1,t}\otimes {\chi_t}^{n-1})(ds_{t,-1}\times dx_{t,-1})\cdot V_n(s_{t1},x_{t1},\varepsilon_{s_{t,-1}x_{t,-1}},\xi_{[t+1,\bar{t}]}). \tag{C.12}$$

We argue that $M_{n31}(\delta)$ can be made arbitrarily small as δ approaches 0^+. Due to Assumption 1 that $\tilde{g}_t(s,x,\tau)$ is continuous in τ at an (s,x)-independent rate, we can make $\tilde{g}_t(s_{t1},x_{t1},\varepsilon_{s_{t,-1}x_{t,-1}})$ for any $(s_{t,-1},x_{t,-1})\in\tilde{U}_{n-1}(\delta)$ arbitrarily close to $\tilde{g}_t(s_{t1},x_{t1},\sigma_t\otimes\chi_t)$ by rendering δ small enough, without respect to (s_{t1},x_{t1}). Due to its countability, we can write $S=\{\bar{s}_1,\bar{s}_2,...\}$. Under known s_{t1}, x_{t1}, $\varepsilon_{s_{t,-1}x_{t,-1}}$, and $\xi_{[t+1,\bar{t}]}$, let us use the simplified notation

$$\gamma_i = \tilde{g}_t(s_{t1},x_{t1},\sigma_t\otimes\chi_t|\{\bar{s}_i\}), \tag{C.13}$$

$$\gamma_i' = \tilde{g}_t(s_{t1},x_{t1},\varepsilon_{s_{t,-1}x_{t,-1}}|\{\bar{s}_i\}), \tag{C.14}$$

and

$$v_i = \Pi_{m=2}^n \int_S \tilde{g}_t(s_{tm},x_{tm},\varepsilon_{s_{t,-m}x_{t,-m}}|ds_{t+1,m})\cdot v_{n,t+1}(s_{t+1,1},\xi_{[t+1,\bar{t}]},\varepsilon_{s_{t+1,-1}},\chi_{[t+1,\bar{t}]}). \tag{C.15}$$

Then, (C.9) can be expressed as

$$V_n(s_{t1},x_{t1},\varepsilon_{s_{t,-1}x_{t,-1}},\xi_{[t+1,\bar{t}]}) = \mid \sum_i \gamma_i\cdot v_i - \sum_i \gamma_i'\cdot v_i \mid . \tag{C.16}$$

Note the $\mid v_i \mid$'s are uniformly bounded, say by $\overline{v}$, due to the boundedness of the $\tilde{f}_{t'}$'s and the finiteness of $\bar{t}$. Let I be the set of i's such that $\gamma_i\geq\gamma_i'$. Then, from (C.16), we have

$$V_n(s_{t1},x_{t1},\varepsilon_{s_{t,-1}x_{t,-1}},\xi_{[t+1,\bar{t}]}) \leq 2\overline{v}\cdot\sum_{i\in I}(\gamma_i-\gamma_i'). \tag{C.17}$$

Let δ be below $\inf_{s\neq s'} d_S(s,s')>0$. But then, $(s_{t,-1},x_{t,-1})\in\tilde{U}_{n-1}(\delta)$ would entail

$$\sum_{i\in I}(\gamma_i-\gamma_i') = \tilde{g}_t(s_{t1},x_{t1},\sigma_t\otimes\chi_t|\{s_i|i\in I\}) - \tilde{g}_t(s_{t1},x_{t1},\varepsilon_{s_{t,-1}x_{t,-1}}|\{s_i|i\in I\}) \\ = \tilde{g}_t(s_{t1},x_{t1},\sigma_t\otimes\chi_t|\{s_i|i\in I\}) - \tilde{g}_t(s_{t1},x_{t1},\varepsilon_{s_{t,-1}x_{t,-1}}|(\{s_i|i\in I\})^\delta) < \delta. \tag{C.18}$$

In view of (C.17), $V_n(s_{t1}, x_{t1}, \varepsilon_{s_{t,-1}x_{t,-1}}, \xi_{[t+1,\bar{t}]})$ with $(s_{t,-1}, x_{t,-1}) \in \tilde{U}_{n-1}(\delta)$ can be made arbitrarily small by decreasing δ at a rate independent of $(s_{t1}, x_{t1}, \xi_{[t+1,\bar{t}]})$. In view of (C.11), we see that $M_{n31}(\delta)$ can be made arbitrarily small by rendering δ small enough.

As noted earlier, the probability $(\hat{\pi}_{n-1,t} \otimes {\chi_t}^{n-1})((S^{n-1} \times X^{n-1}) \setminus \tilde{U}_{n-1}(\delta))$ can be made arbitrarily small at any δ when n is made large enough. But since $V_n(s_{t1}, x_{t1}, \varepsilon_{s_{t,-1}x_{t,-1}}, \xi_{[t+1,\bar{t}]})$ is uniformly bounded, this means that $M_{n12}(\delta)$ can be made arbitrarily small as well.

Hence, all three terms can be made arbitrarily small by letting n be large enough. We have thus completed the induction process. ∎

Proof of Theorem 2: Given (23) for every $t = 1, 2, ..., \bar{t}$ and $\xi_t \in \mathcal{K}(S, X)$, we are to verify (27) for every $t = 1, 2, ..., \bar{t}$, $\epsilon > 0$, large enough n, $s_{t1} \in S$, and $\xi_{[t\bar{t}]} \in (\mathcal{K}(S, X))^{\bar{t}-t+1}$.

First, we show that the one-time formulation of (23) would already imply the futility of any multi-period unilateral deviation. Another way to write the condition is, at $t' = 0$, for any $t = 1, 2, ..., \bar{t} - t'$ and $\xi_{[t,t+t']} \in (\mathcal{K}(S, X))^{t'+1}$,

$$v_t(s_t, (\xi_{[t,t+t'-1]}, \chi_{[t+t',\bar{t}]}), \sigma_t, \chi_{[t\bar{t}]}) \geq v_t(s_t, (\xi_{[t,t+t']}, \chi_{[t+t'+1,\bar{t}]}), \sigma_t, \chi_{[t\bar{t}]}). \tag{C.19}$$

Now suppose (C.19) is true for some $t' = 0, 1, ..., \bar{t} - 1$. We are to show its validity at $t' + 1$. But by (22), for any $t = 1, 2, ..., \bar{t} - t'$,

$$\begin{aligned} &v_t(s_t, (\xi_{[t,t+t']}, \chi_{[t+t'+1,\bar{t}]}), \sigma_t, \chi_{[t\bar{t}]}) - v_t(s_t, (\xi_{[t,t+t'+1]}, \chi_{[t+t'+2,\bar{t}]}), \sigma_t, \chi_{[t\bar{t}]}) \\ &\quad = \textstyle\int_X \xi_t(s_t|dx_t) \cdot \int_S \tilde{g}_t(s_t, x_t, \sigma_t \otimes \chi_t|ds_{t+1}) \times \\ &\qquad \times [v_{t+1}(s_{t+1}, (\xi_{[t+1,t+t']}, \chi_{[t+t'+1,\bar{t}]}), T_t(\chi_t) \circ \sigma_t, \chi_{[t+1,\bar{t}]}) \\ &\qquad\quad - v_{t+1}(s_{t+1}, (\xi_{[t+1,t+t'+1]}, \chi_{[t+t'+2,\bar{t}]}), T_t(\chi_t) \circ \sigma_t, \chi_{[t+1,\bar{t}]})], \end{aligned} \tag{C.20}$$

which, by the induction hypothesis (C.19), is positive. Therefore, (C.19) is true for any $t = 1, 2, ..., \bar{t}$, $t' = 0, 1, ..., \bar{t} - t$, and $\xi_{[t,t+t']} \in (\mathcal{K}(S, X))^{t'+1}$.

Using (C.19) multiple times, we can derive, for any $t = 1, 2, ..., \bar{t}$ and $\xi_{[t\bar{t}]} \in (\mathcal{K}(S, X))^{\bar{t}-t+1}$,

$$\begin{aligned} v_t(s_t, \chi_{[t\bar{t}]}, \sigma_t, \chi_{[t\bar{t}]}) &\geq v_t(s_t, (\xi_t, \chi_{[t+1,\bar{t}]}), \sigma_t, \chi_{[t\bar{t}]}) \geq v_t(s_t, (\xi_{[t,t+1]}, \chi_{[t+2,\bar{t}]}), \sigma_t, \chi_{[t\bar{t}]}) \\ &\geq \cdots \geq v_t(s_t, (\xi_{[t,\bar{t}-1]}, \chi_{\bar{t}}), \sigma_t, \chi_{[t\bar{t}]}) \geq v_t(s_t, \xi_{[t\bar{t}]}, \sigma_t, \chi_{[t\bar{t}]}). \end{aligned} \tag{C.21}$$

In view of (C.21), we would have (27) if for any ϵ and large enough n,

$$\int_{S^{n-1}} \hat{\pi}_{n-1,t}(ds_{t,-1}) \cdot v_{nt}(s_{t1}, \chi_{[t\bar{t}]}, \varepsilon_{s_{t,-1}}, \chi_{[t\bar{t}]}) > v_t(s_{t1}, \chi_{[t\bar{t}]}, \sigma_t, \chi_{[t\bar{t}]}) - \frac{\epsilon}{2}, \tag{C.22}$$

and for any $\xi_{[t\bar{t}]} \in (\mathcal{K}(S, X))^{\bar{t}-t+1}$,

$$\int_{S^{n-1}} \hat{\pi}_{n-1,t}(ds_{t,-1}) \cdot v_{nt}(s_{t1}, \xi_{[t\bar{t}]}, \varepsilon_{s_{t,-1}}, \chi_{[t\bar{t}]}) < v_t(s_{t1}, \xi_{[t\bar{t}]}, \sigma_t, \chi_{[t\bar{t}]}) + \frac{\epsilon}{2}. \tag{C.23}$$

Both (C.22) and (C.23) would be true if

$$\int_{S^{n-1}} \hat{\pi}_{n-1,t}(ds_{t,-1}) \cdot v_{nt}(s_{t1}, \xi_{[t\bar{t}]}, \varepsilon_{s_{t,-1}}, \chi_{[t\bar{t}]}) \longrightarrow_{n\to+\infty} v_t(s_{t1}, \xi_{[t\bar{t}]}, \sigma_t, \chi_{[t\bar{t}]}), \tag{C.24}$$

at an $(s_{t1}, \xi_{[t\bar{t}]})$-independent convergence rate. But this was provided by Proposition 2. ∎

D Proofs of Section 7

Proof of Proposition 3: Since A is discrete, we can denote it by either $\{\bar{a}_1, \bar{a}_2, ...\}$ or $\{\bar{a}_1, ..., \bar{a}_I\}$ for some finite I. We work with the former only, as the latter is similarly treatable.

For any $n \in \mathbb{N}$, define

$$N_n = \{(n_1, n_2, ...) | n_i = 0, 1, ..., n \text{ for each } i = 1, 2, ..., \text{ and } \sum_{i=1}^{+\infty} n_i = n\}. \tag{D.1}$$

For each $(n_1, n_2, ...) \in N_n$, define $A^n_{n_1 n_2 \cdots}$ so that

$$A^n_{n_1 n_2 \cdots} = \{a \in A^n | \varepsilon_a(\{\bar{a}_i\}) = \frac{n_i}{n} \text{ for any } i = 1, 2, ...\}. \tag{D.2}$$

Note that every $A^n_{n_1 n_2 \cdots}$ is symmetric, different $A^n_{n_1 n_2 \cdots}$'s are non-overlapping, and

$$A^n = \bigcup_{(n_1 n_2 \cdots) \in N_n} A^n_{n_1 n_2 \cdots}. \tag{D.3}$$

Due to the above decomposition, each $a \in A^n$ belongs to its own $A^n_{n\cdot\varepsilon_a(\{\bar{a}_1\}), n\cdot\varepsilon_a(\{\bar{a}_2\}), \cdots}$.

For any $(n_1, n_2, \cdots) \in N_n$, the set $A^n_{n_1 n_2 \cdots}$ contains $n!/(\prod_{i=1}^{+\infty} n_i!)$ distinct members of A^n, say $a_1, ..., a_{n!/(\prod_{i=1}^{+\infty} n_i!)}$. In addition, every a_k is of the form ψa_1 for some $\psi \in \Psi_n$. Thus, due to q_n's symmetry, for $k = 1, 2, ..., n!/(\prod_{i=1}^{+\infty} n_i!)$,

$$q_n(\{a_k\}) = \frac{\prod_{i=1}^{+\infty} n_i!}{n!} \cdot q_n(A^n_{n_1 n_2 \cdots}). \tag{D.4}$$

Suppose $n_i \geq 1$ for some $i = 1, 2, ...$ Then, exactly $(n-1)!/((n_i - 1)! \cdot \prod_{j \neq i} n_j!)$ of the a_k's will have $a_{k1} = \bar{a}_i$. Therefore, for any such a_k,

$$q_n((\{\bar{a}_i\} \times A^{n-1}) \cap A^n_{n_1 n_2 \cdots}) = \frac{(n_i - 1)! \cdot \prod_{j \neq i} n_j!}{(n-1)!} \cdot q_n(\{a_k\}) = \frac{n_i}{n} \cdot q_n(A^n_{n_1 n_2 \cdots}), \tag{D.5}$$

where the second equality stems from (D.4). The above left- and right-hand sides are certainly equated as well when $n_i = 0$. Combine (D.3) and (D.5), and we can obtain

$$q_n(\{\bar{a}_i\} \times A^{n-1}) = \sum_{(n_1, n_2, ...) \in N_n} \frac{n_i}{n} \cdot q_n(A^n_{n_1 n_2 \cdots}). \tag{D.6}$$

On the other hand, we have $\min_{i\neq j} d_A(\bar{a}_i, \bar{a}_j) > 0$ due to A's discreteness. Suppose $\epsilon > 0$ is small enough to be strictly below this constant. Then by the nature of the Prohorov metric, $a \in A^n$ would satisfy $\rho_A(\varepsilon_a, p) < \epsilon$ if and only if

$$\sum_{i=1}^{+\infty} \mid \varepsilon_a(\{\bar{a}_i\}) - p(\{\bar{a}_i\}) \mid < 2\epsilon, \tag{D.7}$$

and hence only if

$$\max_{i=1}^{+\infty} \mid \varepsilon_a(\{\bar{a}_i\}) - p(\{\bar{a}_i\}) \mid < \epsilon. \tag{D.8}$$

Since the sequence q_n asymptotically resembles p, for any $\epsilon > 0$ that is strictly below $\min_{i\neq j} d_A(\bar{a}_i, \bar{a}_j) > 0$, we can pick n large enough so that (A.7) and (A.8) in the proof of Lemma 5 are true. Define $N'_n \subseteq N_n$ so that for any $(n_1, n_2, ...) \in N'_n$,

$$\sum_{i=1}^{+\infty} \mid \frac{n_i}{n} - p(\{\bar{a}_i\}) \mid < 2\epsilon, \quad \text{and hence} \quad \max_{i=1}^{+\infty} \mid \frac{n_i}{n} - p(\{\bar{a}_i\}) \mid < \epsilon. \tag{D.9}$$

Due to (D.2), (D.3), and (D.7), we have the following for the A'_n defined in (A.8):

$$A'_n = \bigcup_{(n_1,n_2,...)\in N'_n} A^n_{n_1 n_2 ...}. \tag{D.10}$$

Now for any $i = 1, 2, ...$, we have

$$\begin{aligned} \mid q_n|_A(\{\bar{a}_i\}) - p(\{\bar{a}_i\}) \mid &= \mid q_n(\{\bar{a}_i\} \times A^{n-1}) - p(\{\bar{a}_i\}) \mid \\ &= \mid (\textstyle\sum_{(n_1,n_2,...)\in N'_n} + \sum_{(n_1,n_2,...)\in N_n \setminus N'_n}) n_i \cdot q_n(A^n_{n_1 n_2 ...})/n - p(\{\bar{a}_i\}) \mid \\ &\leq q_n(A'_n) \cdot \mid n_i/n - p(\{\bar{a}_i\}) \mid + q_n(A^n \setminus A'_n) < 2\epsilon. \end{aligned} \tag{D.11}$$

Here, the first equality comes from the definition of marginal probability, the second equality comes from (D.6), the first inequality can be attributed to (D.10), and the last inequality is due to (A.7) and (D.8). Thus, for every $a \in A$, we have $\lim_{n\to+\infty} q_n|_A(\{a\}) = p(\{a\})$. ∎

Proof Proposition 4: For the time being, it does not matter whether $A = \{\bar{a}_1, \bar{a}_2, ...\}$ or $\{\bar{a}_1, ..., \bar{a}_I\}$ for some finite I. The first few steps are the same as those in the proof of Lemma 5. For any $\epsilon > 0$ and n large enough, we can have (A.7) to (A.11) as in that proof.

Fix $i = 1, 2, ...$ with $p(\{\bar{a}_i\}) > 0$. Due to (A.10),

$$A'_n \cap (\{\bar{a}_i\} \times A^{n-1}) \subseteq (A \times A''_{n-1}) \cap (\{\bar{a}_i\} \times A^{n-1}) = \{\bar{a}_i\} \times A''_{n-1}, \tag{D.12}$$

where A'_n is defined in (A.8) and A''_{n-1} is defined in (A.11). Thus,

$$\begin{aligned} q_n(\{\bar{a}_i\} \times A''_{n-1}) &\geq q_n(A'_n \cap (\{\bar{a}_i\} \times A^{n-1})) = q_n((\{\bar{a}_i\} \times A^{n-1}) \setminus (A^n \setminus A'_n)) \\ &\geq q_n(\{\bar{a}_i\} \times A^{n-1}) - q_n(A^n \setminus A'_n) > q_n(\{\bar{a}_i\} \times A^{n-1}) - \epsilon, \end{aligned} \tag{D.13}$$

where the last inequality is due to (A.7).

Since q_n is symmetric, we know from Proposition 3 that, when n is large enough,

$$q_n(\{\bar{a}_i\} \times A^{n-1}) = q_n|_A(\{\bar{a}_i\}) > \frac{p(\{\bar{a}_i\})}{2} > 0. \tag{D.14}$$

Combining (D.13) and (D.14), we can obtain

$$q_{n,A}|_{A^{n-1}}(\bar{a}_i|A''_{n-1}) = \frac{q_n(\{\bar{a}_i\} \times A''_{n-1})}{q_n(\{\bar{a}_i\} \times A^{n-1})} > 1 - \frac{\epsilon}{q_n(\{\bar{a}_i\} \times A^{n-1})} > 1 - \frac{2\epsilon}{p(\{\bar{a}_i\})}. \tag{D.15}$$

With A''_{n-1}'s definition in (A.11), we get $q_{n,A}|_{A^{n-1}}(\bar{a}_i|\cdot)$'s asymptotic resemblance to p^{n-1}. ∎

E Developments in Section 8

Proof of Theorem 4: Let $\epsilon > 0$ be fixed. Given $t = 1, 2, ...$ and $\chi \in \mathcal{K}(S, X)$, we use χ^t to denote $(\chi, \chi, ..., \chi) \in (\mathcal{K}(S, X))^t$. From (38), we know

$$\mid v_n^\infty(s_1, \xi_{[1\infty]}, \varepsilon_{s_{-1}}, \chi^\infty) - v_n^t(s_1, \xi_{[1t]}, \varepsilon_{s_{-1}}, \chi^t) \mid \leq \frac{\bar{\alpha}^t \cdot \bar{f}}{1 - \bar{\alpha}}. \tag{E.1}$$

Hence, when $t \geq \ln(6\bar{f}/(\epsilon \cdot (1 - \bar{\alpha})))/\ln(1/\bar{\alpha}) + 1$,

$$v_n^\infty(s_1, \chi^\infty, \varepsilon_{s_{-1}}, \chi^\infty) > v_n^t(s_1, \chi^t, \varepsilon_{s_{-1}}, \chi^t) - \frac{\epsilon}{6}, \tag{E.2}$$

and

$$v_n^\infty(s_1, \xi_{[1\infty]}, \varepsilon_{s_{-1}}, \chi^\infty) < v_n^t(s_1, \xi_{[1t]}, \varepsilon_{s_{-1}}, \chi^t) + \frac{\epsilon}{6}, \tag{E.3}$$

for every $s_1 \in S$, $s_{-1} \in S^{n-1}$, and $\xi_{[1\infty]} \in (\mathcal{K}(S, X))^\infty$. Therefore, we need merely to select such a large t and show that, when n is large enough,

$$\int_{S^{n-1}} \hat{\pi}_{n-1}(ds_{-1}) \cdot v_n^t(s_1, \chi^t, \varepsilon_{s_{-1}}, \chi^t) \geq \int_{S^{n-1}} \hat{\pi}_{n-1}(ds_{-1}) \cdot v_n^t(s_1, \xi_{[1t]}, \varepsilon_{s_{-1}}, \chi^t) - \frac{2\epsilon}{3}, \tag{E.4}$$

for every $s_1 \in S$ and $\xi_{[1t]} \in (\mathcal{K}(S, X))^t$.

Since (χ, σ) poses as an equilibrium for Γ, we know (36) is true. Another way to write the condition is, at $t' = 0$, for any $\xi_{[1,t'+1]} \in (\mathcal{K}(S, X))^{t'+1}$,

$$v^\infty(s, (\xi_{[1t']}, \chi^\infty), \sigma, \chi^\infty) \geq v^\infty(s, (\xi_{[1,t'+1]}, \chi^\infty), \sigma, \chi^\infty). \tag{E.5}$$

Now suppose (E.5) is true for some $t' = 0, 1, ...$ We are to show its validity at $t' + 1$. By (33), (35), and the uniform convergence of $v^t(s, \xi_{[1t]}, \sigma, \chi^t)$ to $v^\infty(s, \xi_{[1\infty]}, \sigma, \chi^\infty)$, we have

$$\begin{aligned} v^\infty(s, \xi_{[1\infty]}, \sigma, \chi^\infty) = \textstyle\int_X \xi_1(s|dx) \cdot [\tilde{f}(s, x, \sigma \otimes \chi) \\ + \bar{\alpha} \cdot \textstyle\int_S \tilde{g}(s, x, \sigma \otimes \chi|ds') \cdot v^\infty(s', \xi_{[2\infty]}, \sigma, \chi^\infty)]. \end{aligned} \tag{E.6}$$

Therefore,

$$\begin{aligned} & v^{\infty}(s,(\xi_{[1,t'+1]},\chi^{\infty}),\sigma,\chi^{\infty})-v^{\infty}(s,(\xi_{[1,t'+2]},\chi^{\infty}),\sigma,\chi^{\infty}) \\ & \quad =\int_X \xi_1(s|dx)\cdot\int_S \tilde{g}(s,x,\sigma\otimes\chi|ds')\times \\ & \qquad \times[v^{\infty}(s',(\xi_{[2,t'+1]},\chi^{\infty}),\sigma,\chi^{\infty})-v^{\infty}(s',(\xi_{[2,t'+2]},\chi^{\infty}),\sigma,\chi^{\infty})], \end{aligned} \tag{E.7}$$

which, by the induction hypothesis (E.5), is positive. Therefore, (E.5) is true for $t'=0,1,....$

By using (E.5) multiple times, we can derive that, for any $\xi_{[1t]}\in(\mathcal{K}(S,X))^t$,

$$\begin{aligned} v^{\infty}(s,\chi^{\infty},\sigma,\chi^{\infty}) & \geq v^{\infty}(s,(\xi_1,\chi^{\infty}),\sigma,\chi^{\infty})\geq v^{\infty}(s,(\xi_{[12]},\chi^{\infty}),\sigma,\chi^{\infty}) \\ & \geq\cdots\geq v^{\infty}(s,(\xi_{[1,t-1]},\chi^{\infty}),\sigma_t,\chi^{\infty})\geq v^{\infty}(s,(\xi_{[1t]},\chi^{\infty}),\sigma,\chi^{\infty}). \end{aligned} \tag{E.8}$$

Also, we know from (34) that

$$\mid v^{\infty}(s,\zeta_{[1\infty]},\sigma,\chi^{\infty})-v^{t}(s,\zeta_{[1t]},\sigma,\chi^{t})\mid\leq\frac{\bar{\alpha}^t\cdot\bar{f}}{1-\bar{\alpha}}, \tag{E.9}$$

regardless of the $\zeta_{[1\infty]}\in(\mathcal{K}(S,X))^{\infty}$ chosen. However, (E.8) and (E.9) would together lead to

$$v^{t}(s,\chi^{t},\sigma,\chi^{t})-v^{t}(s,\xi_{[1t]},\sigma,\chi^{t})\geq-\frac{2\bar{\alpha}^{t-1}\cdot\bar{f}}{1-\bar{\alpha}}\geq-\frac{\epsilon}{3}, \tag{E.10}$$

for any $s\in S$ and $\xi_{[1t]}\in(\mathcal{K}(S,X))^t$.

In the presence of Assumptions 1 and 2 for the corresponding t-period games, Proposition 2 applies. Plus, it has been hypothesized that the sequence $\hat{\pi}_{n-1}$ asymptotically resembles the sequence σ^{n-1}. Therefore, for n large enough,

$$\int_{S^{n-1}}\hat{\pi}_{n-1}(ds_{-1})\cdot v_n^t(s_1,\chi^t,\varepsilon_{s_{-1}},\chi^t)>v^t(s_1,\chi^t,\sigma,\chi^t)-\frac{\epsilon}{6}, \tag{E.11}$$

regardless of the choice on $s_1\in S$, and

$$\int_{S^{n-1}}\hat{\pi}_{n-1}(ds_{-1})\cdot v_n^t(s_1,\xi_{[1t]},\varepsilon_{s_{-1}},\chi^t)<v^t(s_1,\xi_{[1t]},\sigma,\chi^t)+\frac{\epsilon}{6}, \tag{E.12}$$

regardless of the choices on $s_1\in S$ and $\xi_{[1t]}\in(\mathcal{K}(S,X))^t$. Put (E.10) to (E.12) together, and we would obtain (E.4). ∎

For something akin to the second example in Section 7, we need to consider the following invariant equation involving $\pi_n\in\mathcal{P}(S^n)$, which is inspirable from its finite-t version (28):

$$\pi_n=\pi_n\odot\chi^n\odot\tilde{g}^n. \tag{E.13}$$

Suppose (E.13) has a solution that asymptotically resembles σ^n, then we can let $\hat{\pi}_{n-1}=\pi_n|_{S^{n-1}}$. By Lemma 5, this choice would satisfy the condition in Theorem 4. Its meaning is

also clear—let players update their estimates on other players' states most precisely without using their own state information.

When the state space S is finite, we again have an extended version much like Theorem 3. If we succeed in finding a satisfactory π_n, we would be able to make the third choice of letting each $\hat{\pi}_{n-1}(s_1|\cdot)$ in the extended version be the conditional probability $\pi_{n,S}|_{S^{n-1}}(s_1|\cdot)$. Propositions 3 and 4 would then lead to the satisfaction of the corresponding condition in the extended version. The third choice here again means that players update other players' states in the most accurate Bayesian fashion.

The above second and third choices are premised on the following conjecture.

Conjecture 1 *Suppose $\chi \in \mathcal{K}(S,X)$, $\tilde{g} \in \mathcal{G}(S,X)$ enjoys the continuity of $\tilde{g}(s,x,\tau)$ in τ at an (s,x)-independent rate, and $\sigma \in \mathcal{P}(S)$ is an solution to the invariant equation $\sigma = \sigma \odot \chi \odot \tilde{g}(\cdot,\cdot,\sigma \otimes \chi)$ as defined by (32) and (35). Then, there would exist a sequence π_n so that for each $n \in \mathbb{N}$, π_n as a member of $\mathcal{P}(S^n)$ satisfies the invariant equation $\pi_n = \pi_n \odot \chi^n \odot \tilde{g}^n$ as indicated by (E.13), and yet the sequence asymptotically resembles the sequence σ^n.*

To tackle this conjecture, one may be tempted to show that (i) iteratively applying $\sigma_{t+1} = \sigma_t \odot \chi \odot \tilde{g}(\cdot,\cdot,\sigma_t \otimes \chi)$ leads to the convergence of σ_t to an invariant σ, (ii) iteratively applying $\pi_{n,t+1} = \pi_{nt} \odot \chi^n \odot \tilde{g}^n$ leads to the convergence of π_{nt} to an invariant π_n for each n, and (iii) these convergence results along with the asymptotic resemblance of each π_{nt} to ${\sigma_t}^n$ would lead to that of π_n to σ^n. So far, (i) and (ii) still elude us. On the other hand, something slightly weaker than (iii) can be achieved.

Proposition 5 *Let A be a separable metric space, and p_i for $i \in \mathbb{N}$ and p be members of $\mathcal{P}(A)$. Also, for each $n \in \mathbb{N}$, let q_{ni} for $i \in \mathbb{N}$ and q_n be members of $\mathcal{P}(A^n)$. Suppose p_i converges to p, q_{ni} converges to q_n for each $n \in \mathbb{N}$, and q_{ni} asymptotically resembles ${p_i}^n$. Then, in either situation (a) where the convergence of q_{ni} to q_n is at an n-independent rate or situation (b) where the asymptotic resemblance of q_{ni} to ${p_i}^n$ is at an i-independent rate, the sequence q_n would asymptotically resemble the sequence p^n.*

Proof of Proposition 5: Let $\epsilon > 0$ be given. Since p_i converges to p, we have

$$\rho_A(p,p_i) < \frac{\epsilon}{2}, \tag{E.14}$$

as long as i is large enough.

Suppose situation (a) is true. By the equi-n convergence of q_{ni} to q_n, we can pick i large enough to ensure both (E.14) and for any $n \in \mathbb{N}$,

$$q_n((A'_n)^{\epsilon/4}) > q_{ni}(A') - \frac{\epsilon}{2}, \qquad \forall A'_n \in \mathcal{B}(A^n). \tag{E.15}$$

At such a fixed $i \in \mathbb{N}$, due to the asymptotic resemblance of q_{ni} to ${p_i}^n$, we can let n be large enough so that

$$q_{ni}(\{a \in A^n | \rho_A(\varepsilon_a, p_i) < \frac{\epsilon}{4}\}) > 1 - \frac{\epsilon}{2}. \tag{E.16}$$

Suppose situation (b) is true. Due to the equi-i asymptotic resemblance of q_{ni} to ${p_i}^n$, we can pick n large enough to ensure (E.16) for any $i \in \mathbb{N}$. By the convergence of q_{ni} to q_n, we can then pick i large enough to ensure (E.14), as well as (E.15) for the current $n \in \mathbb{N}$.

Either way, without loss of generality, we can suppose $d_{A^n}(a, a') \geq \max_{m=1}^{n} d_A(a_m, a'_m)$. Then, due to Lemma 1,

$$(\{a \in A^n | \rho_A(\varepsilon_a, p_i) < \frac{\epsilon}{4}\})^{\epsilon/4} \subseteq \{a \in A^n | \rho_A(\varepsilon_a, p_i) < \frac{\epsilon}{2}\}. \tag{E.17}$$

Now we can deduce that

$$\begin{aligned} q_n(\{a \in A^n | \rho_A(\varepsilon_a, p) < \epsilon\}) &> q_n(\{a \in A^n | \rho_A(\varepsilon_a, p_i) < \epsilon/2\}) \\ &> q_n((\{a \in A^n | \rho_A(\varepsilon_a, p_i) < \epsilon/4\})^{\epsilon/4}) \\ &> q_{ni}(\{a \in A^n | \rho_A(\varepsilon_a, p_i) < \epsilon/4\}) - \epsilon/2 > 1 - \epsilon, \end{aligned} \tag{E.18}$$

where the first inequality is due to (E.14), the second inequality is due to (E.17), the third inequality is due to (E.15), and the last inequality is due to (E.16). Therefore, the sequence q_n asymptotically resembles the sequence p^n. ∎

Like Propositions 3 and 4, Proposition 5 also helps to bolster the legitimacy of the asymptotic resemblance concept.

F Developments in Section 9

F.1 The Transient Case

By the discreteness of S, every $\chi_t(s|X')$ is automatically continuous and hence measurable in s, and hence $\mathcal{K}(S, X)$ is not only a member of $(\mathcal{P}(X))^S$, but also the latter itself. Denote

the space $(\mathcal{P}(S))^{\bar{t}-1}$ by $\mathcal{S}$ and the space $((\mathcal{P}(X))^S)^{\bar{t}} = (\mathcal{K}(S,X))^{\bar{t}}$ by $\mathcal{X}$. Let $\mathcal{U} = \mathcal{S} \times \mathcal{X}$. Define a correspondence $H : \mathcal{U} \Rightarrow \mathcal{U}$, so that for any $\sigma_{[2\bar{t}]} \in \mathcal{S}$ and $\chi_{[1\bar{t}]} \in \mathcal{X}$,

$$H(\sigma_{[2\bar{t}]}, \chi_{[1\bar{t}]}) = H^S(\sigma_{[2\bar{t}]}, \chi_{[1\bar{t}]}) \times H^X(\sigma_{[2\bar{t}]}, \chi_{[1\bar{t}]}), \tag{F.1}$$

where

$$H^S(\sigma_{[2\bar{t}]}, \chi_{[1\bar{t}]}) = \{\sigma'_{[2\bar{t}]} \in \mathcal{S} | \sigma'_t = T_{t-1}(\chi_{t-1}) \circ \sigma_{t-1},\ \forall t = 2, 3, ..., \bar{t}\}, \tag{F.2}$$

and

$$H^X(\sigma_{[2\bar{t}]}, \chi_{[1\bar{t}]}) = \{\chi'_{[1\bar{t}]} \in \mathcal{X} | \chi'_t(s_t | \tilde{X}_t(s_t, \sigma_t, \chi_{[t\bar{t}]})) = 1,\ \forall t = 1, 2, ..., \bar{t}, s_t \in S\}. \tag{F.3}$$

A fixed point $(\sigma_{[2\bar{t}]}, \chi_{[1\bar{t}]})$ for H would provide a Markov equilibrium $\chi_{[1\bar{t}]}$ for $\Gamma(\sigma_1)$ in the sense of (42), with $\sigma_{[2\bar{t}]}$ supplying the deterministic pre-action environment pathway from period 2 to $\bar{t}$ that is generated from all players adopting policy $\chi_{[1\bar{t}]}$. We are to use Kakutani-Fan-Glicksberg fixed point theorem to prove the existence of a fixed point for H. But first let us work out a couple of useful continuity results.

Proposition 6 *(i) $\sigma \otimes \chi$ is continuous in both $\sigma \in \mathcal{P}(S)$ and $\chi \in (\mathcal{P}(X))^S$.*
When $g \in \mathcal{G}(S,X)$ satisfies that $g(s,x,\tau)$ is continuous in τ at an (s,x)-independent rate,
(ii) $\sigma \odot \chi \odot g(\cdot,\cdot,\sigma \otimes \chi)$ is continuous in both $\sigma \in \mathcal{P}(S)$ and $\chi \in (\mathcal{P}(X))^S$.

Proof of Proposition 6: We first prove (i) by showing that, for any two sequences σ_m and χ_m that converge to σ and χ, respectively, the sequence $\sigma_m \otimes \chi_m$ would converge to $\sigma \otimes \chi$. In the following, we omit detailed reasonings behind some of the steps, as they have appeared in the proof of Proposition 1.

Fix some $\epsilon \in (0,1)$. We can identify some I of its points $\bar{s}_1, \bar{s}_2, ..., \bar{s}_I$, so that (B.1) is true. For convenience, let $\bar{S}' = \{\bar{s}_1, \bar{s}_2, ..., \bar{s}_I\}$ and $\bar{S}'' = S \setminus \bar{S}'$. It is known that the distance $d_S(\bar{S}', \bar{S}'') = \inf_{s' \in \bar{S}', s'' \in \bar{S}''} d_S(s', s'') > 0$. For $i, j = 1, 2, ..., I$, use d_{ij} for $d_S(\bar{s}_i, \bar{s}_j)$ and σ_i for $\sigma(\{\bar{s}_i\})$. Again, define δ through (B.2), whose strict positivity is guaranteed.

As σ_m and χ_m converge to σ and χ, respectively, for large enough m, we have

$$\rho_S(\sigma, \sigma_m) < \delta, \tag{F.4}$$

and

$$\rho_X(\chi(\bar{s}_i), \chi_m(\bar{s}_i)) < \delta, \qquad \forall i = 1, 2, ..., I. \tag{F.5}$$

Together with the fact that $\delta \leq d_S(\bar{S}', \bar{S}'') \wedge (\min_{i \neq j} d_{ij})$, (F.4) would result with

$$\sigma_i - \delta < \sigma_m(\{\bar{s}_i\}) < \sigma_i + \delta. \tag{F.6}$$

Meanwhile, (F.5) would lead to

$$\chi_m(\bar{s}_i|X') < \chi(\bar{s}_i|(X')^\delta) + \delta, \qquad \forall i = 1, 2, ..., I. \tag{F.7}$$

Any $U' \in \mathcal{B}(S \times X)$ still enjoys the decomposition provided in (B.12), that $U' = (\bigcup_{i=1}^{I} \{\bar{s}_i\} \times X'_i) \bigcup U''$, where $X'_i \in \mathcal{B}(X)$ for $i = 1, 2, ..., I$, while U'' is such that $s'' \in \bar{S}''$ for any $(s'', x'') \in U''$. This would result in the same (B.13). On the other hand, from the right half of (F.6) and (F.7),

$$\begin{aligned}(\sigma_m \otimes \chi_m)(\{\bar{s}_i\} \times X'_i) = \sigma_m(\{\bar{s}_i\}) \cdot \chi_m(\bar{s}_i|X'_i) < (\sigma_i + \delta) \cdot [\chi(\bar{s}_i|(X'_i)^\delta) + \delta] \\ \leq (\sigma \otimes \chi)(\{\bar{s}_i\} \times (X'_i)^\delta) + 2\delta + \delta^2 < (\sigma \otimes \chi)(\{\bar{s}_i\} \times (X'_i)^\delta) + 3\delta,\end{aligned} \tag{F.8}$$

where the last inequality is due to our choice that $\delta \leq \epsilon/I < 1$. Meanwhile,

$$\begin{aligned}(\sigma_m \otimes \chi_m)(U'') \leq (\sigma_m \otimes \chi_m)(\bar{S}'' \times X) = \sigma_m(\bar{S}'') = 1 - \textstyle\sum_{i=1}^{I} \sigma_m(\{\bar{s}_i\}) \\ < 1 - \textstyle\sum_{i=1}^{I} \sigma_i + I\delta < \epsilon + I\delta,\end{aligned} \tag{F.9}$$

where the second-to-last inequality is due to the left half of (F.6) and the last one is due to (B.1). By combining (B.12), (B.13), (F.8), and (F.9), we can obtain

$$(\sigma_m \otimes \chi_m)(U') < (\sigma \otimes \chi)((U')^\delta) + \epsilon + 4I\delta. \tag{F.10}$$

Thus,

$$\rho_{S \times X}(\sigma_m \otimes \chi_m, \sigma \otimes \chi) < \epsilon + 4I\delta \leq 5\epsilon. \tag{F.11}$$

Since (F.11) is to occur at any m that is large enough, we see that (i) is true.

We then prove (ii). Again, suppose two sequences σ_m and χ_m converge to σ and χ, respectively. From (i), we know $\sigma_m \otimes \chi_m$ converges to $\sigma \otimes \chi$ too. According to (87) of Yang [33], for any m,

$$\rho_X(\sigma_m \odot \chi_m, \sigma \odot \chi) = \rho_X((\sigma_m \otimes \chi_m)|_X, (\sigma \otimes \chi)|_X) \leq \rho_{S \times X}(\sigma_m \otimes \chi_m, \sigma \otimes \chi). \tag{F.12}$$

Hence, there is also the convergence of $\sigma_m \odot \chi_m$ to $\sigma \odot \chi$.

On the other hand, the discrete property of $S \times X$ means $g(\cdot, \cdot, \tau)$ is a member of $(\mathcal{P}(S))^{S \times X}$ for any fixed $\tau \in \mathcal{P}(S \times X)$. Now (i) and the fact that $g(s, x, \tau)$ is continuous in τ at an (s, x)-independent rate would together mean that, the sequence $g(\cdot, \cdot, \sigma_m \otimes \chi_m)$ in $(\mathcal{P}(S))^{S \times X}$ converges to $g(\cdot, \cdot, \sigma \otimes \chi)$.

Let us use the convergence of $\sigma_m \odot \chi_m$ to $\sigma \odot \chi$ under proper substitutions. As $S \times X$ has been noted to be discrete, we can treat it as S in the convergence result. Also, let us treat $\sigma_m \otimes \chi_m$ as σ_m, $\sigma \otimes \chi$ as σ, S as X, $g(\cdot, \cdot, \sigma_m \otimes \chi_m)$ as χ_m, and $g(\cdot, \cdot, \sigma \otimes \chi)$ as χ.

From (i) on the convergence of $\sigma_m \otimes \chi_m$ to $\sigma \otimes \chi$, now viewed as that of σ_m to σ, as well as the convergence of $g(\cdot,\cdot,\sigma_m \otimes \chi_m)$ to $g(\cdot,\cdot,\sigma \otimes \chi)$, now viewed as that of χ_m to χ, we can conclude that $(\sigma_m \otimes \chi_m) \odot g(\cdot,\cdot,\sigma_m \otimes \chi_m) = \sigma_m \odot \chi_m \odot g(\cdot,\cdot,\sigma_m \otimes \chi_m)$ would converge to $(\sigma \otimes \chi) \odot g(\cdot,\cdot,\sigma \otimes \chi) = \sigma \odot \chi \odot g(\cdot,\cdot,\sigma \otimes \chi)$. Thus, (ii) is true as well. ∎

Proposition 7 *For each $t = 1, 2, ..., \bar{t}+1$, the value $v_t(s_t, \xi_{[t\bar{t}]}, \sigma_t, \chi_{[t\bar{t}]})$ defined in (22) is continuous in $\sigma_t \in \mathcal{S}$ and $\chi_{[t\bar{t}]} \in \mathcal{X}$ at an $(s_t, \xi_{[t\bar{t}]})$-independent rate.*

Proof of Proposition 7: We use induction on t. By (21), our claim is certainly true for $t = \bar{t}+1$. Suppose for some $t = \bar{t}, \bar{t}-1, ..., 1$, the function $v_{t+1}(s_{t+1}, \xi_{[t+1,\bar{t}]}, \sigma_{t+1}, \chi_{[t+1,\bar{t}]})$ is continuous in σ_{t+1} and $\chi_{[t+1,\bar{t}]}$ at a rate independent of s_{t+1} and $\xi_{[t+1,\bar{t}]}$.

Now we prove the continuity in σ_t and $\chi_{[t\bar{t}]}$ at time t. From (22), we have

$$\sup_{s_t \in S, \xi_{[t\bar{t}]} \in ((\mathcal{P}(X))^S)^{\bar{t}-t+1}} \mid v_t(s_t, \xi_{[t\bar{t}]}, \sigma_t, \chi_{[t\bar{t}]}) - v_t(s_t, \xi_{[t\bar{t}]}, \sigma'_t, \chi'_{[t\bar{t}]}) \mid \leq M_1 + M_2 + M_3, \quad \text{(F.13)}$$

where

$$M_1 = \sup_{(s_t,x_t) \in S \times X} \mid \tilde{f}_t(s_t, x_t, \sigma_t \otimes \chi_t) - \tilde{f}_t(s_t, x_t, \sigma'_t \otimes \chi'_t) \mid, \quad \text{(F.14)}$$

$$\begin{aligned} M_2 = \sup_{(s_t,x_t) \in S \times X,\, \xi_{[t+1,\bar{t}]} \in ((\mathcal{P}(X))^S)^{\bar{t}-t}} \mid & [\textstyle\int_S \tilde{g}_t(s_t, x_t, \sigma_t \otimes \chi_t | ds_{t+1}) \\ & - \textstyle\int_S \tilde{g}_t(s_t, x_t, \sigma'_t \otimes \chi'_t | ds_{t+1})] \cdot v_{t+1}(s_{t+1}, \xi_{[t+1,\bar{t}]}, T_t(\chi_t) \circ \sigma_t, \chi_{[t+1,\bar{t}]}) \mid, \end{aligned} \quad \text{(F.15)}$$

and

$$\begin{aligned} M_3 = & \sup_{(s_t,x_t) \in S \times X,\, \xi_{[t+1,\bar{t}]} \in ((\mathcal{P}(X))^S)^{\bar{t}-t}} \textstyle\int_S \tilde{g}_t(s_t, x_t, \sigma'_t \otimes \chi'_t | ds_{t+1}) \times \\ & \times \mid v_{t+1}(s_{t+1}, \xi_{[t+1,\bar{t}]}, T_t(\chi_t) \circ \sigma_t, \chi_{[t+1,\bar{t}]}) - v_{t+1}(s_{t+1}, \xi_{[t+1,\bar{t}]}, T_t(\chi'_t) \circ \sigma'_t, \chi'_{[t+1,\bar{t}]}) \mid . \end{aligned} \quad \text{(F.16)}$$

By part (i) of Proposition 6, $\sigma'_t \otimes \chi'_t$ can be made arbitrarily close to $\sigma_t \otimes \chi_t$ by letting (σ'_t, χ'_t) be close enough to (σ_t, χ_t). Then due to Assumption 2, M_1 can be made arbitrarily small by doing the same.

Again, suppose $S = \{\bar{s}_1, \bar{s}_2, ...\}$. We use the simplified notation that

$$\gamma_i^{(\prime)}(s_t, x_t) = g_t(s_t, x_t, \sigma_t^{(\prime)} \otimes \chi_t^{(\prime)} | \{\bar{s}_i\}), \quad \text{(F.17)}$$

and

$$v_i(\xi_{[t+1,\bar{t}]}) = v_{t+1}(\bar{s}_i, \xi_{[t+1,\bar{t}]}, T_t(\chi_t) \circ \sigma_t, \chi_{[t+1,\bar{t}]}). \quad \text{(F.18)}$$

Then, (F.15) can be expressed as M_2 equaling

$$\sup_{(s_t,x_t) \in S \times X,\, \xi_{[t+1,\bar{t}]} \in ((\mathcal{P}(X))^S)^{\bar{t}-t}} \mid \sum_i \gamma_i(s_t, x_t) \cdot v_i(\xi_{[t+1,\bar{t}]}) - \sum_i \gamma'_i(s_t, x_t) \cdot v_i(\xi_{[t+1,\bar{t}]}) \mid . \quad \text{(F.19)}$$

Let $I(s_t, x_t)$ be the set of i's that induce $\gamma_i(s_t, x_t) \geq \gamma'_i(s_t, x_t)$. Note the $\mid v_i(\xi_{[t+1,\bar{t}]}) \mid$'s are bounded, say by $\overline{v}$, due to the boundedness of the $\tilde{f}_{t'}$'s and the finiteness of $\bar{t}$. Then, (F.19) would lead to

$$M_2 \leq 2\overline{v} \cdot \sup_{(s_t,x_t)\in S\times X} \sum_{i\in I(s_t,x_t)} (\gamma_i(s_t, x_t) - \gamma'_i(s_t, x_t)). \tag{F.20}$$

For δ below $\inf_{s\neq s'} d_S(s, s')$, the event $\rho_S(\tilde{g}_t(s_t, x_t, \sigma_t \otimes \chi_t), \tilde{g}_t(s_t, x_t, \sigma'_t \otimes \chi'_t)) < \delta$ would trigger

$$\sum_{i\in I(s_t,x_t)} (\gamma_i(s_t, x_t) - \gamma'_i(s_t, x_t)) < \delta, \tag{F.21}$$

for every $(s_t, x_t) \in S \times X$; consult (C.18) in the proof of Proposition 2. But due to Assumption 1, the convergence of $\sigma'_t \otimes \chi'_t$ to $\sigma_t \otimes \chi_t$ means that we can make $\tilde{g}_t(s_t, x_t, \sigma'_t \otimes \chi'_t)$ arbitrarily close to $\tilde{g}_t(s_t, x_t, \sigma_t \otimes \chi_t)$, at a rate that is independent of (s_t, x_t). Hence, by (F.20), M_2 can be made arbitrarily small by letting (σ'_t, χ'_t) get close enough to (σ_t, χ_t).

From (F.16), we can get

$$\begin{aligned} M_3 \leq \sup_{s_{t+1}\in S,\, \xi_{[t+1,\bar{t}]}\in((\mathcal{P}(X))^S)^{\bar{t}-t}} \mid & v_{t+1}(s_{t+1}, \xi_{[t+1,\bar{t}]}, T_t(\chi_t) \circ \sigma_t, \chi_{[t+1,\bar{t}]}) \\ & - v_{t+1}(s_{t+1}, \xi_{[t+1,\bar{t}]}, T_t(\chi'_t) \circ \sigma'_t, \chi'_{[t+1,\bar{t}]}) \mid . \end{aligned} \tag{F.22}$$

By part (ii) of Proposition 6, $T_t(\chi'_t) \circ \sigma'_t = \sigma'_t \odot \chi'_t \odot \tilde{g}_t(\cdot, \cdot, \sigma'_t \otimes \chi'_t)$ can be made arbitrarily close to $T_t(\chi_t) \circ \sigma_t = \sigma_t \odot \chi_t \odot \tilde{g}_t(\cdot, \cdot, \sigma_t \otimes \chi_t)$ by letting (σ'_t, χ'_t) be close enough to (σ_t, χ_t). By the induction hypothesis, M_3 can be made arbitrarily small by doing the same.

We have thus completed the induction process. ∎

Here comes the conditional-equilibrium existence result for the transient case.

Theorem 5 *The correspondence H allows for a fixed point $(\sigma_{[2\bar{t}]}, \chi_{[1\bar{t}]})$, which supplies the game $\Gamma(\sigma_1)$ with an conditional equilibrium $\chi_{[1\bar{t}]}$.*

Proof of Theorem 5: Due to S's discreteness, $\mathcal{P}(S)$ is the simplex in $\mathbb{R}^{|S|}$, whether $\mid S \mid$ be finite or infinite, and hence is compact; the same applies to $\mathcal{P}(X)$. Thus, $\mathcal{U}$ is a compact subset of the vector space $\mathbb{R}^{|S|^{\bar{t}-1}+|X|^{|S|\cdot\bar{t}}}$, understood as $\mathbb{R}^\infty$ if either S or X is infinite.

For any finite-dimensional $\mathbb{R}^k$, we can take the norm $||\cdot||$ so that $|| r ||= \sum_{l=1}^k \mid r_l \mid /k$ for each $r = (r_l | l = 1, ..., k) \in \mathbb{R}^k$, whereas for the infinite-dimensional $\mathbb{R}^\infty$, we can let $|| r ||= \sum_{l=1}^{+\infty} \mid r_l \mid /2^l$ for each $r = (r_l | l = 1, 2, ...) \in \mathbb{R}^\infty$. A norm thus defined would provide the same convergence as does the weak convergence under Prohorov metric. Since the convex combination of two probabilities is still a probability, $\mathcal{U}$ is also convex.

For any $(\sigma_{[2\bar{t}]}, \chi_{[1\bar{t}]}) \in \mathcal{U}$, the set $H(\sigma_{[2\bar{t}]}, \chi_{[1\bar{t}]})$ is certainly non-empty, for we can construct some $(\sigma'_{[2\bar{t}]}, \chi'_{[1\bar{t}]})$ belonging to it. First, for $t = 2, 3, ..., \bar{t}$, we simply let $\sigma'_t = T_{t-1}(\chi_{t-1}) \circ \sigma_{t-1}$. Then, for $t = 1, 2, ..., \bar{t}$ and $s \in S$, let $\chi'_t(s)$ be any measure that assigns its full weight to the set of x's that attain the maximum value $\sup_{y \in X} v_t(s, (\delta_y, \chi_{[t+1,\bar{t}]}), \sigma_t, \chi_{[t\bar{t}]})$.

Now we show that $H^S : \mathcal{U} \Rightarrow \mathcal{S}$ and $H^X : \mathcal{U} \Rightarrow \mathcal{X}$ are closed- and convex-valued, as well as upper hemi-continuous. These would lead to the same properties for H. According to (F.2), each $H^S(\sigma_{[2\bar{t}]}, \chi_{[1\bar{t}]})$ contains exactly one point, and hence is automatically closed and convex. For the upper hemi-continuity property, we need only to show that the value contained in $H^S(\sigma_{[2\bar{t}]}, \chi_{[1\bar{t}]})$ moves continuously with both $\sigma_{[2\bar{t}]}$ and $\chi_{[1\bar{t}]}$. But this has been guaranteed by part (ii) of Proposition 6.

According to (F.3), each $H^X(\sigma_{[2\bar{t}]}, \chi_{[1\bar{t}]})$ is a set of probability vectors, with each component probability assigning the full measure to a particular measurable set. This set of probability vectors is certainly convex. To show that it is closed, suppose $\chi'_{m,[1\bar{t}]}$ for $m = 1, 2, ...$ form a sequence in $H^X(\sigma_{[2\bar{t}]}, \chi_{[1\bar{t}]})$ that converges to a given $\chi'_{[1\bar{t}]}$. We are to show that

$$\chi'_{[1\bar{t}]} \in H^X(\sigma_{[2\bar{t}]}, \chi_{[1\bar{t}]}). \tag{F.23}$$

Now for any $t = 1, 2, ..., \bar{t}$, $s \in S$, and $\epsilon > 0$, as long as m is large enough,

$$\chi'_t(s|(\tilde{X}_t(s, \sigma_{[2\bar{t}]}, \chi_{[1\bar{t}]}))^\epsilon) \geq \chi'_{mt}(s|\tilde{X}_t(s, \sigma_{[2\bar{t}]}, \chi_{[1\bar{t}]})) - \epsilon = 1 - \epsilon. \tag{F.24}$$

Due to the arbitrariness of ϵ, this means $\chi'_t(s|\tilde{X}_t(s, \sigma_{[2\bar{t}]}, \chi_{[1\bar{t}]})) = 1$, and hence (F.23) is true.

We now show that H^X is upper hemi-continuous. Let $\sigma_{m,[2\bar{t}]}$ be a sequence in $\mathcal{S}$ that converges to a given $\sigma_{[2\bar{t}]}$, $\chi_{m,[1\bar{t}]}$ a sequence in $\mathcal{X}$ that converges to a given $\chi_{[1\bar{t}]}$, and $\chi'_{m,[1\bar{t}]}$ another sequence in $\mathcal{X}$ that converges to a given $\chi'_{[1\bar{t}]}$. Suppose for each $m = 1, 2, ...,$

$$\chi'_{m,[1\bar{t}]} \in H^X(\sigma_{m,[2\bar{t}]}, \chi_{m,[1\bar{t}]}), \tag{F.25}$$

we are to show that

$$\chi'_{[1\bar{t}]} \in H^X(\sigma_{[2\bar{t}]}, \chi_{[1\bar{t}]}). \tag{F.26}$$

By (F.3), we see that (F.25) for each m indicates that, for each $t = 1, 2, ..., \bar{t}$ and $s \in S$,

$$\chi'_{mt}(s|\tilde{X}_t(s, \sigma_{mt}, \chi_{m,[t\bar{t}]})) = 1; \tag{F.27}$$

whereas, (F.26) boils down to that, for each $t = 1, 2, ..., \bar{t}$ and $s \in S$,

$$\chi'_t(s|\tilde{X}_t(s, \sigma_t, \chi_{[t\bar{t}]})) = 1. \tag{F.28}$$

We fix some t and s. Let $\epsilon > 0$ be small enough, so that there is no need to distinguish between $(X')^\epsilon$ and X' for any $X' \subseteq X$. Now since χ'_{mt} converges to χ'_t, for m large enough,

$$\chi'_t(s|\tilde{X}_t(s,\sigma_t,\chi_{[t\bar{t}]})) = \chi'_t(s|(\tilde{X}_t(s,\sigma_t,\chi_{[t\bar{t}]}))^\epsilon) \geq \chi'_{mt}(s|\tilde{X}_t(s,\sigma_t,\chi_{[t\bar{t}]})) - \epsilon. \tag{F.29}$$

For the time being, suppose $\tilde{X}_t(s,\sigma_t,\chi_{[t\bar{t}]})$ is known to be upper hemi-continuous in $(\sigma_t,\chi_{[t\bar{t}]})$. By also noting the hypothesis on the convergence of σ_{mt} to σ_t and that of $\chi_{m,[t\bar{t}]}$ to $\chi_{[t\bar{t}]}$, we can obtain, for m large enough,

$$\tilde{X}_t(s,\sigma_{mt},\chi_{m,[t\bar{t}]}) \subseteq (\tilde{X}_t(s,\sigma_t,\chi_{[t\bar{t}]}))^\epsilon = \tilde{X}_t(s,\sigma_t,\chi_{[t\bar{t}]}). \tag{F.30}$$

Thus, for m large enough,

$$\chi'_t(s|\tilde{X}_t(s,\sigma_t,\chi_{[t\bar{t}]})) \geq \chi'_{mt}(s|\tilde{X}_t(s,\sigma_t,\chi_{[t\bar{t}]})) - \epsilon \geq \chi'_{mt}(s|\tilde{X}_t(s,\sigma_{mt},\chi_{m,[t\bar{t}]})) - \epsilon, \tag{F.31}$$

which, according to (F.27), is above $1-\epsilon$. In view of the arbitrariness of ϵ, we can achieve (F.28).

We now come back to the upper hemi-continuity of $\tilde{X}_t(s,\cdot)$ as a correspondence from $\mathcal{P}(S) \times ((\mathcal{P}(X))^S)^{\bar{t}-t+1}$ to X. Suppose σ_{mt} converges to σ_t, $\chi_{m,[t\bar{t}]}$ converges to $\chi_{[t\bar{t}]}$, and x_m converges to x. For every $m = 1, 2, ...$, suppose $x_m \in \tilde{X}_t(s,\sigma_{mt},\chi_{m,[t\bar{t}]})$, which, by (43), means

$$v_t(s,(\delta_{x_m},\chi_{m,[t+1,\bar{t}]}),\sigma_{mt},\chi_{m,[t\bar{t}]}) \geq v_t(s,(\delta_y,\chi_{m,[t+1,\bar{t}]}),\sigma_{mt},\chi_{m,[t\bar{t}]}), \qquad \forall y \in X. \tag{F.32}$$

By X's discreteness, x_m would be x for sufficiently large m. This, combined with Proposition 7 and the hypothesis on the convergence of σ_{mt} to σ_t and that of $\chi_{m,[t\bar{t}]}$ to $\chi_{[t\bar{t}]}$, would entail that, for any $\epsilon > 0$, as long as m is large enough,

$$\begin{aligned} v_t(s,(\delta_x,\chi_{[t+1,\bar{t}]}),\sigma_t,\chi_{[t\bar{t}]}) &\geq v_t(s,(\delta_x,\chi_{m,[t+1,\bar{t}]}),\sigma_{mt},\chi_{m,[t\bar{t}]}) - \epsilon \\ &= v_t(s,(\delta_{x_m},\chi_{m,[t+1,\bar{t}]}),\sigma_{mt},\chi_{m,[t\bar{t}]}) - \epsilon \\ &\geq v_t(s,(\delta_y,\chi_{m,[t+1,\bar{t}]}),\sigma_{mt},\chi_{m,[t\bar{t}]}) - \epsilon \geq v_t(s,(\delta_y,\chi_{[t+1,\bar{t}]}),\sigma_t,\chi_{[t\bar{t}]}) - 2\epsilon, \end{aligned} \tag{F.33}$$

for any $y \in X$. Since ϵ can be arbitrarily small, we see from (43) that $x \in \tilde{X}_t(s,\sigma_t,\chi_{[t\bar{t}]})$. Thus we have the upper hemi-continuity of $\tilde{X}_t(s,\cdot)$.

In summary, H is a non-empty, closed- and convex-valued, as well as upper hemi-continuous correspondence on the compact and convex subset $\mathcal{U}$ that is embedded in a normed linear topological space. We can therefore apply the Kakutani-Fan-Glicksberg fixed point theorem to verify that H has a fixed point. ∎

F.2 The Stationary Case

Denote the space $\mathcal{P}(S)$ by $\mathcal{S}_\infty$ and the space $(\mathcal{P}(X))^S$ by $\mathcal{X}_\infty$. Let $\mathcal{U}_\infty = \mathcal{S}_\infty \times \mathcal{X}_\infty$. Define a correspondence $H_\infty : \mathcal{U}_\infty \Rightarrow \mathcal{U}_\infty$, so that for any $\sigma \in \mathcal{S}_\infty$ and $\chi \in \mathcal{X}_\infty$,

$$H_\infty(\sigma, \chi) = H_\infty^S(\sigma, \chi) \times H_\infty^X(\sigma, \chi), \tag{F.34}$$

where

$$H_\infty^S(\sigma, \chi) = \{\sigma' \in \mathcal{S}_\infty | \sigma' = T(\chi) \circ \sigma\}, \tag{F.35}$$

and

$$H_\infty^X(\sigma, \chi) = \{\chi' \in \mathcal{X}_\infty | \chi'(s|\tilde{X}_\infty(s, \sigma, \chi)) = 1, \ \forall s \in S\}. \tag{F.36}$$

A fixed point (σ, χ) for H_∞ would provide a stationary Markov equilibrium χ for the stationary nonatomic game Γ^∞ in the sense of (44), with σ supplying the invariant deterministic environment that is generated from all players adopting policy χ. To show that such an equilibrium exists, we first need the following consequence of Proposition 7.

Proposition 8 *The value $v^\infty(s, \xi_{[1\infty]}, \sigma, \chi^\infty)$ defined in (33) is continuous in $\sigma \in \mathcal{S}_\infty$ and $\chi \in \mathcal{X}_\infty$ at an $(s, \xi_{[1\infty]})$-independent rate.*

Proof of Proposition 8: From (34), we see that

$$\mid v^\infty(s, \xi_{[1\infty]}, \sigma, \chi^\infty) - v^t(s, \xi_{[1t]}, \sigma, \chi^t) \mid \leq \frac{\bar{\alpha}^t \cdot \overline{f}}{1 - \bar{\alpha}}. \tag{F.37}$$

Thus, for any $\epsilon > 0$, by fixing at a large enough t, we can ensure

$$\mid v^\infty(s, \xi_{[1\infty]}, \sigma'', (\chi'')^\infty) - v^t(s, \xi_{[1t]}, \sigma'', (\chi'')^t) \mid < \frac{\epsilon}{3}, \tag{F.38}$$

for any s, $\xi_{[1\infty]}$, σ'', and χ''. At the same time, Proposition 7 means that, for (σ', χ') close enough to any given (σ, χ), we can guarantee

$$\mid v^t(s, \xi_{[1t]}, \sigma, \chi^t) - v^t(s, \xi_{[1t]}, \sigma', (\chi')^t) \mid < \frac{\epsilon}{3}, \tag{F.39}$$

for any s and $\xi_{[1t]}$. Then,

$$\begin{aligned} \mid v^\infty(s, \xi_{[1\infty]}, \sigma, \chi^\infty) - v^\infty(s, \xi_{[1\infty]}, \sigma', (\chi')^\infty) \mid \leq & \mid v^\infty(s, \xi_{[1\infty]}, \sigma, \chi^\infty) \\ & - v^t(s, \xi_{[1t]}, \sigma, \chi^t) \mid + \mid v^t(s, \xi_{[1t]}, \sigma, \chi^t) - v^t(s, \xi_{[1t]}, \sigma', (\chi')^t) \mid \\ & + \mid v^t(s, \xi_{[1t]}, \sigma', (\chi')^t) - v^\infty(s, \xi_{[1\infty]}, \sigma', (\chi')^\infty) \mid < \epsilon. \end{aligned} \tag{F.40}$$

Thus, $v^{\infty}(s, \xi_{[1\infty]}, \sigma, \chi^{\infty})$ is continuous in (σ, χ) at an $(s, \xi_{[1\infty]})$-independent rate. ∎

We can then have the desired conditional-equilibrium existence result by using the Kakutani-Fan-Glicksberg fixed point theorem.

Theorem 6 *The correspondence H_{∞} allows for a fixed point (σ, χ), which supplies the game Γ^{∞} with an equilibrium χ.*

Proof of Theorem 6: Due to the discreteness of S and X, $\mathcal{U}_{\infty}$ is a compact subset of the vector space $\mathbb{R}^{|S|+|X|^{|S|}}$, understood as $\mathbb{R}^{\infty}$ if either S or X is infinite. Regardless of whether the space is finite- or infinite-dimensional, we can take the norm adopted in the proof of Theorem 5. Since the convex combination of two probabilities is still a probability, $\mathcal{U}_{\infty}$ is convex.

Using virtually the same corresponding arguments in the proof of Theorem 5, we can show that $H_{\infty}(\sigma, \chi)$ at any $(\sigma, \chi) \in \mathcal{U}_{\infty}$ is non-empty, closed, and convex. We separate the upper hemi-continuity of H_{∞} into that for H_{∞}^{S} and that for H_{∞}^{X}.

The upper hemi-continuity of H_{∞}^{S} again comes from Proposition 6. Furthermore, we can use almost the same arguments from (F.32) to (F.33), this time relying on Proposition 8 instead of Proposition 7, to show that $\tilde{X}_{\infty}(s, \cdot)$ as a correspondence from $\mathcal{P}(S) \times (\mathcal{P}(X))^{S}$ to X is upper hemi-continuous. Then, using almost the same arguments from (F.25) to (F.31), we can verify that H_{∞}^{X} is upper hemi-continuous.

With all these properties, we can apply the Kakutani-Fan-Glicksberg fixed point theorem to verify that H_{∞} has a fixed point. ∎

References

[1] Adlakha, S. and B. Johari. 2013. Mean Field Equilibrium in Dynamic Games with Complementarities. *Operations Research*, **61**, pp. 971-989.

[2] Al-Najjar, N.I. and R. Smorodinsky. 2001. Large Nonanonymous Repeated Games. *Games and Economic Behavior*, **37**, pp. 26-39.

[3] Al-Najjar, N.I. 2008. Large Games and the Law of Large Numbers. *Games and Economic Behavior*, **64**, pp. 1-34.

[4] Balder, E.J. 2002. A Unifying Pair of Cournot-Nash Equilibrium Existence Results. *Journal of Economic Theory*, **102**, pp. 437-470.

[5] Balder, E.J. 2008. Comments on Purification in Continuous Games. *International Journal of Game Theory*, **37**, pp. 73-92.

[6] Bergin, J. and D. Bernhardt. 1995. Anonymous Sequential Games: Existence and Characterization of Equilibria. *Economic Theory*, **5**, pp. 461-489.

[7] Bodoh-Creed, A.L. 2012. Approximation of Large Dynamic Games. Working Paper, Cornell University.

[8] Carmona, G. 2004. Nash Equilibria of Games with a Continuum of Players. Working Paper, Universidade Nova de Lisboa.

[9] Duffie, D., J. Geanakoplos, A. Mas-Colell, and A. McLennan. 1994. Stationary Markov Equilibria. *Econometrica*, **62**, pp. 745-781.

[10] Dvoretzky, A., J. Kiefer, and J. Wolfolwitz. 1956. Asymptotic Minimax Character of the Sample Distribution Function and of the Classical Multinomial Estimator. *Annals of Mathematical Statistics*, **27**, pp. 642-669.

[11] Ethier, S.N. and T.G. Kurtz. 1986. *Markov Processes: Characterization and Convergence*. John Wiley & Sons, New York.

[12] Green, E.J. 1980. Non-cooperative Price Taking in Large Dynamic Markets. *Journal of Economic Theory*. **22**, pp. 155-182.

[13] Green, E.J. 1984. Continuum and Finite-player Noncooperative Models of Competition. *Econometrica*, **52**, pp. 975-993.

[14] Hopenhayn, H.A. 1992. Entry, Exit, and Firm Dynamics in Long Run Equilibrium. *Econometrica*, **60**, pp. 1127-1150.

[15] Housman, D. 1988. Infinite Player Noncooperative Games and the Continuity of the Nash Equilibrium Correspondence. *Mathematics of Operations Research*, **13**, pp. 488-496.

[16] Jovanovic, B. and R.W. Rosenthal. 1988. Anonymous Sequential Games. *Journal of Mathematical Economics*, **17**, pp. 77-88.

[17] Kalai, E. 2004. Large Robust Games. *Econometrica*, **72**, pp. 1631-1665.

[18] Khan, M.A., K.P. Rath, and Y.N. Sun. 1997. On the Existence of Pure Strategy Equilibrium in Games with a Continuum of Players. *Journal of Economic Theory*, **76**, pp. 13-46.

[19] Khan, M.A., K.P. Rath, Y.N. Sun, and H. Yu. 2013. Large Games with a Bio-social Typology. *Journal of Economic Theory*, **148**, 1122-1149.

[20] Khan, M.A. and Y.N. Sun. 1990. On a Reformulation of Cournot-Nash Equilibria. *Journal of Mathematical Analysis and Applications*, **146**, pp. 442-460.

[21] Khan, M.A. and Y.N. Sun. 1995. Pure Strategies in Games with Private Information. *Journal of Mathematical Economics*, **24**, pp. 633-653.

[22] Khan, M.A. and Y.N. Sun. 1999. Non-cooperative Games on Hyperfinite Loeb Space. *Journal of Mathematical Economics*, **31**, pp. 455-492.

[23] Khan, M.A. and Y.N. Sun. 2002. Non-cooperative Games with Many Players. In R.J. Aumann and S. Hart (eds.), *Handbook of Game Theory with Economic Applications, Volume 3*, Elsevier Science, Amsterdam, pp. 1761-1808.

[24] Mas-Colell, A. 1984. On a Theorem of Schmeidler. *Journal of Mathematical Economics*, **13**, pp. 201-206.

[25] Mertens, J.F. and T. Parthasarathy. 1987. Equalibria for Discounted Stochastic Games. CORE Discussion Paper No. 8750.

[26] Parthasarathy, K.R. 2005. *Probability Measures on Metric Spaces*. AMS Chelsea Publishing, Providence, Rhode Island.

[27] Sabourian, H. 1990. Anonymous Repeated Games with a Large Number of Players and Random Outcomes. *Journal of Economic Theory*, **51**, pp. 92-110.

[28] Schmeidler, D. 1973. Equilibrium Points of Nonatomic Games. *Journal of Statistical Physics*, **7**, pp. 295-300.

[29] Shapley, L.S. 1953. Stochastic Games. *Proceedings of the National Academy of Sciences*, **39**, pp. 1095-1100.

[30] Solan, E. 1998. Discounted Stochastic Games. *Mathematics of Operations Research*, **23**, pp. 1010-1021.

[31] Weintraub, G.Y., C. L. Benkard, and B. van Roy. 2008. Markov Perfect Industry Dynamics with Many Firms. *Econometrica*, **76**, pp. 1375-1411.

[32] Weintraub, G.Y., C. L. Benkard, and B. van Roy. 2011. Industry Dynamics: Foundations for Models with an Infinite Number of Firms. *Journal of Economic Theory*, **146**, pp. 504-527.

[33] Yang, J. 2011. Asymptotic Interpretations for Equilibria of Nonatomic Games. *Journal of Mathematical Economics*, **47**, pp. 491-499.

[34] Yang, J. 2015. Analysis of Markovian Competitive Situations using Nonatomic Games—the Shock-driven Case. Working Paper, Rutgers University.

[35] Yang, J. and Y. Xia. 2013. A Nonatomic-game Approach to Dynamic Pricing under Competition. *Production and Operations Management*, **22**, pp. 88-103.

[36] Yu, H. and Z. Zhang. 2007. Pure Strategy Equilibria in Games with Countable Actions. *Journal of Mathematical Econonomics*, **43**, pp. 192-200.